

## User Guide

## **Amazon Batch**





## **Table of Contents**

What is Amazon Batch?	1
Are you a first-time Amazon Batch user?	2
Related services	2
Accessing Amazon Batch	3
Components of Amazon Batch	3
Compute environment	4
Job queues	4
Job definitions	4
Jobs	4
Scheduling policy	5
Consumable resources	5
Setting up Amazon Batch	6
Create IAM account and administrative user	
Sign up for an Amazon Web Services account	6
Secure IAM users	7
Create IAM roles	7
Create a key pair	8
Create a VPC	10
Create a security group	11
Install the Amazon CLI	
Getting started tutorials	. 14
Getting started with Amazon EC2 using the Wizard	. 14
Overview	. 14
Prerequisites	. 15
Step 1: Create a compute environment	. 15
Step 2: Create a job queue	. 16
Step 3: Create a job definition	16
Step 4: Create a job	. 17
Step 5: Review and create	. 17
Step 6: View the Job's output	. 17
Step 7: Clean up your tutorial resources	. 18
Additional resources	. 18
Getting started with Fargate orchestration using the Wizard	. 19
Overview	. 19

Prerequisites	20
Step 1: Create a compute environment	20
Step 2: Create a job queue	21
Step 3: Create a job definition	21
Step 4: Create a job	22
Step 5: Review and create	22
Step 6: View the Job's output	22
Step 7: Clean up your tutorial resources	23
Additional resources	23
Getting started with Amazon Batch and Fargate using the Amazon CLI	24
Prerequisites	24
Create an IAM execution role	25
Create a compute environment	27
Create a job queue	28
Create a job definition	28
Submit and monitor a job	29
View job output	31
Clean up resources	31
Going to production	32
Next steps	33
Getting started with Amazon EKS	34
Overview	35
Prerequisites	36
Step 1: Create your Amazon EKS cluster for Amazon Batch	37
Step 2: Prepare your Amazon EKS cluster for Amazon Batch	37
Step 3: Create an Amazon EKS compute environment	41
Step 4: Create a job queue and attach the compute environment	
Step 5: Create a job definition	44
Step 6: Submit a job	45
Step 7: View the Job's output	45
Step 8: (Optional) Submit a job with overrides	46
Step 9: Clean up your tutorial resources	47
Additional resources	47
Getting started with Amazon Batch on Amazon EKS Private Clusters	48
Overview	35
Prerequisites	50

Step 1: Create your EKS cluster for Amazon Batch	51
Step 2: Prepare your EKS cluster for Amazon Batch	53
Step 3: Create an Amazon EKS compute environment	56
Step 4: Create a job queue and attach the compute environment	58
Step 5: Create an Amazon ECR with pull through cache	59
Step 6: Register a job definition	60
Step 7: Submit a job to run	61
Step 8: View the Job's output	61
Step 9: (Optional) Submit a job with overrides	61
Step 10: Clean up your tutorial resources	62
Additional resources	47
Troubleshooting	63
The Amazon Batch widget dashboard	65
Add the Single job queue widget	65
How to add the CloudWatch Container Insights widget	66
Add the Job logs widget	66
Compute environments for Amazon Batch	
Managed compute environments	68
Consideration when creating multi-node parallel jobs	71
Unmanaged compute environments	
Create a compute environment	72
Tutorial: Create a managed compute environment using Fargate resources	73
Tutorial: Create a managed compute environment using Amazon EC2 resources	75
Tutorial: Create an unmanaged compute environment using Amazon EC2 resources	81
Tutorial: Create a managed compute environment using Amazon EKS resources	82
Resource: Compute environment template	86
Compute resource AMIs	
Compute resource AMI specification	
Tutorial: Create a compute resource AMI	91
Use a GPU workload AMI	95
Amazon Linux deprecation	101
Use Amazon EC2 launch templates	101
Default and override launch templates	103
Amazon EC2 user data in launch templates	104
Reference: Launch template examples	105
EC2 configurations	108

Instance type allocation strategies	109
Memory management	110
Reserve system memory	111
Tutorial: View compute resource memory	112
Memory and vCPU considerations for Amazon Batch on Amazon EKS	112
Updating compute environments	118
Updating Amazon Fargate compute environments	118
Updating the AMI ID	121
Fargate compute environments	122
When to use Fargate	123
Job definitions on Fargate	123
Job queues on Fargate	126
Compute environments on Fargate	126
Amazon EKS compute environments	127
Amazon EKS	129
Default AMI	129
Mixed AMI environments	130
Supported Kubernetes versions	131
Update the Kubernetes version of the compute environment	132
Shared responsibility of the Kubernetes nodes	132
Run a DaemonSet on Amazon Batch managed nodes	133
Customize Amazon EKS launch templates	134
How to upgrade from EKS AL2 to EKS AL2023	138
Job queues	140
Create a job queue	140
Tutorial: Create an Amazon EC2 job queue	140
Tutorial: Create a Fargate job queue	142
Tutorial: Create an Amazon EKS job queue	143
Reference: Job queue template	144
View job queue status	145
View job queue information	145
Fair-share scheduling policies	146
Use share identifiers	147
Use scheduling policies	
Use fair-share scheduling	148
Tutorial: Create a scheduling policy	149

Reference: Scheduling policy template	. 151
Resource-aware scheduling	. 151
Create consumable resources	152
Specify resources for a job	. 153
Check resource usage	155
Update in-use resource quantities	. 156
Find jobs requiring a consumable resource	157
Delete a consumable resource	. 158
Job definitions	160
Create a single-node job definition	160
Tutorial: Create a single-node job definition on Amazon EC2 resources	. 161
Tutorial: Create a single-node job definition on Fargate resources	. 166
Tutorial: Create a single-node job definition on Amazon EKS resources	. 172
Tutorial: Create a single-node job definition with multiple containers on Amazon EC2	
resources	. 177
Create a multi-node parallel job definition	. 183
Tutorial: Create a multi-node parallel job definition on Amazon EC2 resources	184
Reference: Job definition template using ContainerProperties	. 191
Reference: Job definition parameters for ContainerProperties	. 199
Create job definitions using EcsProperties	. 242
ContainerProperties versus EcsProperties job definitions	
General changes to the Amazon Batch APIs	243
Multi-container job definitions for Amazon ECS	
Multi-container job definitions for Amazon EKS	245
Reference: Amazon Batch job scenarios using EcsProperties	. 246
Use the awslogs log driver	252
awslogs log driver options in the Amazon Batch JobDefiniton data type	253
Specify a log configuration in your job definition	. 255
Specify sensitive data	257
Use Secrets Manager	257
Use Systems Manager Parameter Store	265
Private registry authentication for jobs	. 269
Required IAM permissions for private registry authentication	. 271
Tutorial: Create a secret for private registry authentication	271
Amazon EFS volumes	. 272
Amazon EFS volume considerations	. 273

	Use Amazon EFS access points	274
	Specify an Amazon EFS file system in your job definition	275
	Reference: Job definition examples	278
	Reference: Environment variables	278
	Reference: Parameter substitution	279
	Reference: Test GPU functionality	280
	Reference: Multi-node parallel job	281
Jo	bs	283
	Tutorial: submit a job	283
	Job states	286
	Job environment variables	289
	Automated job retries	291
	Job dependencies	292
	Job timeouts	293
	Amazon EKS jobs	294
	Tutorial: Map a running job to a pod and a node	295
	Tutorial: Map a running pod back to its job	296
	Multi-node parallel jobs	297
	Environment variables	298
	Node groups	299
	Job lifecycle	300
	Compute environment considerations	300
	Multi-node parallel jobs on Amazon EKS	301
	Running MNP jobs	302
	Create an Amazon EKS MNP job definition	304
	Submit an Amazon EKS MNP job	306
	Override an Amazon EKS MNP job definition	306
	Array jobs	307
	Reference: Example of an array job workflow	
	Tutorial: Using array job index	313
	Run GPU jobs	319
	Tutorial: Create a GPU-based Kubernetes cluster on Amazon EKS	324
	Tutorial: Create an Amazon EKS GPU job definition	326
	Tutorial: Run a GPU job in your Amazon EKS cluster	326
	Search Amazon Batch jobs in a job queue	327
	Tutorial: Search and filter Amazon Batch jobs	327

	Tutorial: (Beta) Search and filter Amazon Batch jobs	329
	Networking modes for Amazon Batch jobs	330
	Tutorial: Send Amazon Batch job logs to CloudWatch Logs	331
	Tutorial: Review Amazon Batch job information	332
S	ecurity in Amazon Batch	334
	Identity and Access Management	335
	Audience	335
	Authenticating with identities	336
	Managing access using policies	339
	How Amazon Batch works with IAM	341
	Identity-based policy examples	347
	Amazon managed policies	349
	IAM policies, roles, and permissions	364
	IAM policy structure	365
	Resource: Example policies	368
	Resource: Amazon Batch managed policy	376
	Amazon Batch IAM execution role	378
	Supported resource-level permissions	379
	Tutorial: Create the IAM execution role	379
	Tutorial: Check the IAM execution role	380
	Use service-linked roles	381
	Amazon ECS instance role	388
	Amazon EC2 spot fleet role	392
	EventBridge IAM role	395
	Create a virtual private cloud	397
	Create a VPC	397
	Next steps	397
	VPC endpoints	398
	Considerations	398
	Create an interface endpoint	400
	Create an endpoint policy	400
	Compliance validation	401
	Infrastructure security	402
	Cross-service confused deputy prevention	403
	Example: Role for accessing only one compute environment	404
	Example: Role for accessing multiple compute environments	404

CloudTrail	405
Amazon Batch information in CloudTrail	406
Reference: Understanding Amazon Batch log file entries	406
Troubleshoot Amazon Batch IAM	408
I am not authorized to perform an action in Amazon Batch	409
I am not authorized to perform iam:PassRole	409
I want to allow people outside of my Amazon account to access my Amazon Batch	
resources	410
Amazon Step Functions	411
Tutorial: View state machine details	411
Tutorial: Edit a state machine	412
Tutorial: Run a state machine	412
Amazon EventBridge	413
Amazon Batch events	414
Resource: Job state change events	414
Resource: Job queue blocked events	416
Tutorial: Use Amazon user notifications with Amazon Batch	418
Amazon Batch jobs as EventBridge targets	418
Tutorial: Create a scheduled job	419
Tutorial: Create a rule with an event pattern	422
Tutorial: Pass input transformer	424
Tutorial: Listen for Amazon Batch job events	
Prerequisites	427
Tutorial: Create the Lambda function	427
Tutorial: Register the event rule	428
Tutorial: Test your configuration	430
Tutorial: Sending Amazon Simple Notification Service alerts for failed job events	431
Prerequisites	431
Tutorial: Create and subscribe to an Amazon SNS topic	431
Tutorial: Register an event rule	432
Tutorial: Test your rule	434
Alternate rule: Batch job queue blocked	434
Elastic Fabric Adapter	435
Monitor Amazon Batch	438
CloudWatch Logs	438
Tutorial: Add a CloudWatch Logs IAM policy	439

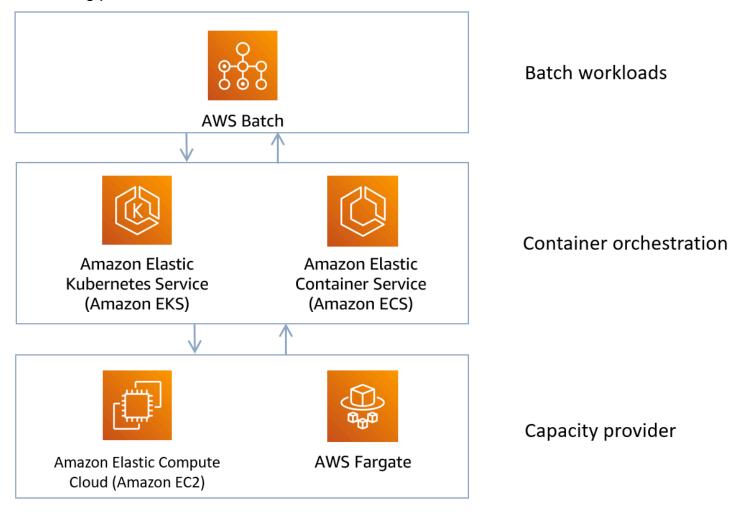
Install and configure the CloudWatch agent	441
Tutorial: View CloudWatch Logs	441
CloudWatch Container Insights	443
Tutorial: Turn on Container Insights	444
Use CloudWatch Logs to monitor Amazon Batch on Amazon EKS jobs	445
Prerequisites	445
Install the add-on	445
Tag your resources	447
Tag basics	447
Tag your resources	448
Tag restrictions	449
Tutorial: Manage tags using the console	450
Add tags on an individual resource on creation	450
Add and delete tags on an individual resource	
Manage tags using the CLI or API	450
Best practices	453
When to use Amazon Batch	453
Checklist to run at scale	
Optimize containers and AMIs	455
Choose the right compute environment resource	456
Amazon EC2 On-Demand or Amazon EC2 Spot	456
Use Amazon EC2 Spot best practices for Amazon Batch	
Common errors and troubleshooting	459
Troubleshooting	462
Amazon Batch	463
INVALID compute environment	463
Jobs stuck in a RUNNABLE status	466
Spot Instances not tagged on creation	471
Spot Instances not scaling down	
Can't retrieve Secrets Manager secrets	473
Can't override job definition resource requirements	473
Error message when you update the desiredvCpus setting	475
Amazon Batch on Amazon EKS	
INVALID compute environment	475
Amazon Batch on Amazon EKS job is stuck in RUNNABLE status	479
Amazon Batch on Amazon EKS job is stuck in STARTING status	480

Verify that the aws-auth ConfigMap is configured correctly	481
RBAC permissions or bindings aren't configured properly	482
Resource: Service quotas	484
Document history	485

## What is Amazon Batch?

Amazon Batch helps you to run batch computing workloads on the Amazon Web Services Cloud. Batch computing is a common way for developers, scientists, and engineers to access large amounts of compute resources. Amazon Batch removes the undifferentiated heavy lifting of configuring and managing the required infrastructure, similar to traditional batch computing software. This service can efficiently provision resources in response to jobs submitted in order to eliminate capacity constraints, reduce compute costs, and deliver results quickly.

As a fully managed service, Amazon Batch helps you to run batch computing workloads of any scale. Amazon Batch automatically provisions compute resources and optimizes the workload distribution based on the quantity and scale of the workloads. With Amazon Batch, there's no need to install or manage batch computing software, so you can focus your time on analyzing results and solving problems.



1

Amazon Batch provides all of the necessary functionality to run high-scale, compute-intensive workloads on top of Amazon managed container orchestration services, Amazon ECS and Amazon EKS. Amazon Batch is able to scale compute capacity on Amazon EC2 instances and Fargate resources.

Amazon Batch provides a fully managed service for batch workloads, and delivers the operational capabilities to optimize these types of workloads for throughput, speed, resource efficiency, and cost.

#### **Topics**

- Are you a first-time Amazon Batch user?
- Related services
- Accessing Amazon Batch
- Components of Amazon Batch

## Are you a first-time Amazon Batch user?

If you are a first-time user of Amazon Batch, we recommend that you begin by reading the following sections:

- · Components of Amazon Batch
- Create IAM account and administrative user
- Setting up Amazon Batch
- Getting started with Amazon Batch tutorials

### **Related services**

Amazon Batch is a fully managed batch computing service that plans, schedules, and runs your containerized batch ML, simulation, and analytics workloads across the full range of Amazon compute offerings, such as Amazon ECS, Amazon EKS, Amazon Fargate, and Spot or On-Demand Instances. For more information about each managed compute service, see:

- Amazon EC2 User Guide
- Amazon Fargate Developer Guide
- Amazon EKS User Guide

## **Accessing Amazon Batch**

You can access Amazon Batch using the following:

#### **Amazon Batch console**

The web interface where you create and manage resources.

#### **Amazon Command Line Interface**

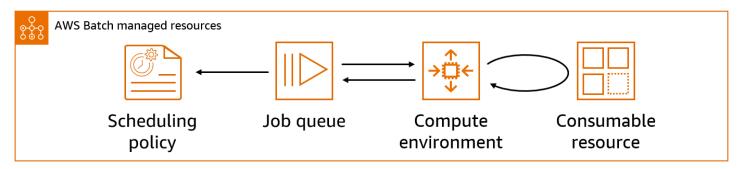
Interact with Amazon Web Services services using commands in your command line shell. The Amazon Command Line Interface is supported on Windows, macOS, and Linux. For more information about the Amazon CLI, see <a href="Amazon Command Line Interface User Guide">Amazon CLI Command Reference</a>. You can find the Amazon Batch commands in the Amazon CLI Command Reference.

#### **Amazon SDKs**

If you prefer to build applications using language-specific APIs instead of submitting a request over HTTP or HTTPS, use the libraries, sample code, tutorials, and other resources provided by Amazon. These libraries provide basic functions that automate tasks, such as cryptographically signing your requests, retrying requests, and handling error responses. These functions make it more efficient for you to get started. For more information, see <u>Tools to Build on Amazon</u>.

## **Components of Amazon Batch**

Amazon Batch simplifies running batch jobs across multiple Availability Zones within a Region. You can create Amazon Batch compute environments within a new or existing VPC. After a compute environment is up and associated with a job queue, you can define job definitions that specify which Docker container images to run your jobs. Container images are stored in and pulled from container registries, which may exist within or outside of your Amazon infrastructure.



Accessing Amazon Batch 3

## **Compute environment**

A compute environment is a set of managed or unmanaged compute resources that are used to run jobs. With managed compute environments, you can specify desired compute type (Fargate or EC2) at several levels of detail. You can set up compute environments that use a particular type of EC2 instance, a particular model such as c5.2xlarge or m5.10xlarge. Or, you can choose only to specify that you want to use the newest instance types. You can also specify the minimum, desired, and maximum number of vCPUs for the environment, along with the amount that you're willing to pay for a Spot Instance as a percentage of the On-Demand Instance price and a target set of VPC subnets. Amazon Batch efficiently launches, manages, and terminates compute types as needed. You can also manage your own compute environments. As such, you're responsible for setting up and scaling the instances in an Amazon ECS cluster that Amazon Batch creates for you. For more information, see Compute environments for Amazon Batch.

## Job queues

When you submit an Amazon Batch job, you submit it to a particular job queue, where the job resides until it's scheduled onto a compute environment. You associate one or more compute environments with a job queue. You can also assign priority values for these compute environments and even across job queues themselves. For example, you can have a high priority queue that you submit time-sensitive jobs to, and a low priority queue for jobs that can run anytime when compute resources are cheaper. For more information, see Job queues.

### **Job definitions**

A job definition specifies how jobs are to be run. You can think of a job definition as a blueprint for the resources in your job. You can supply your job with an IAM role to provide access to other Amazon resources. You also specify both memory and CPU requirements. The job definition can also control container properties, environment variables, and mount points for persistent storage. Many of the specifications in a job definition can be overridden by specifying new values when submitting individual Jobs. For more information, see <u>Job definitions</u>

### **Jobs**

A unit of work (such as a shell script, a Linux executable, or a Docker container image) that you submit to Amazon Batch. It has a name, and runs as a containerized application on Amazon Fargate or Amazon EC2 resources in your compute environment, using parameters that you specify in a job definition. Jobs can reference other jobs by name or by ID, and can be dependent on the successful

Compute environment 4

completion of other jobs or the availability of <u>resources</u> you specify. For more information, see Jobs.

## **Scheduling policy**

You can use scheduling policies to configure how compute resources in a job queue are allocated between users or workloads. Using fair-share scheduling policies, you can assign different share identifiers to workloads or users. The Amazon Batch job scheduler defaults to a first-in, first-out (FIFO) strategy. For more information, see <u>Fair-share scheduling policies</u>.

#### Consumable resources

A consumable resource is a resource that is needed to run your jobs, such as a 3rd party license token, database access bandwidth, the need to throttle calls to a third-party API, and so on. You specify the consumable resources which are needed for a job to run, and Batch takes these resource dependencies into account when it schedules a job. You can reduce the under-utilization of compute resources by allocating only the jobs that have all the required resources available. For more information, see <a href="Resource-aware scheduling">Resource-aware scheduling</a>.

Scheduling policy 5

## **Setting up Amazon Batch**

If you've already signed up for Amazon Web Services (Amazon) and are using Amazon Elastic Compute Cloud (Amazon EC2) or Amazon Elastic Container Service (Amazon ECS), you can soon use Amazon Batch. The setup process for these services is similar. This is because Amazon Batch uses Amazon ECS container instances in its compute environments. To use the Amazon CLI with Amazon Batch, you must use a version of the Amazon CLI that supports the latest Amazon Batch features. If you don't see support for an Amazon Batch feature in the Amazon CLI, upgrade to the latest version. For more information, see http://www.amazonaws.cn/cli/.



#### Note

Because Amazon Batch uses components of Amazon EC2, you use the Amazon EC2 console for many of these steps.

Complete the following tasks to get set up for Amazon Batch.

#### **Topics**

- Create IAM account and administrative user
- Create IAM roles for your compute environments and container instances
- Create a key pair for your instances
- Create a VPC
- Create a security group
- Install the Amazon CLI

## Create IAM account and administrative user

To get started, you need to create an Amazon account and a single user that is typically granted administrative rights. To accomplish this, complete the following tutorials:

## Sign up for an Amazon Web Services account

If you do not have an Amazon Web Services account, use the following procedure to create one.

#### To sign up for Amazon Web Services

- 1. Open http://www.amazonaws.cn/ and choose **Sign Up**.
- 2. Follow the on-screen instructions.

Amazon sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <a href="http://www.amazonaws.cn/">http://www.amazonaws.cn/</a> and choosing **My Account**.

#### Secure IAM users

After you sign up for an Amazon Web Services account, safeguard your administrative user by turning on multi-factor authentication (MFA). For instructions, see <a href="Enable a virtual MFA device for an IAM user (console)">Enable a virtual MFA device for an IAM user (console)</a> in the IAM User Guide.

To give other users access to your Amazon Web Services account resources, create IAM users. To secure your IAM users, turn on MFA and only give the IAM users the permissions needed to perform their tasks.

For more information about creating and securing IAM users, see the following topics in the *IAM User Guide*:

- Creating an IAM user in your Amazon Web Services account
- Access management for Amazon resources
- Example IAM identity-based policies

## Create IAM roles for your compute environments and container instances

Your Amazon Batch compute environments and container instances require Amazon Web Services account credentials to make calls to other Amazon APIs on your behalf. Create an Amazon Identity and Access Management role that provides these credentials to your compute environments and container instances, then associate that role with your compute environments.

Secure IAM users 7



#### Note

To verify that your Amazon Web Services account has the required permissions, see Initial IAM service set up for your account.

The Amazon Batch compute environment and container instance roles are automatically created for you in the console first-run experience. So, if you intend to use the Amazon Batch console, you can move ahead to the next section. If you plan to use the Amazon CLI instead, complete the procedures in Use service-linked roles for Amazon Batch, Amazon ECS instance role, and Tutorial: Create the IAM execution role before creating your first compute environment.

## Create a key pair for your instances

Amazon uses public-key cryptography to secure the login information for your instance. A Linux instance, such as an Amazon Batch compute environment container instance, has no password to use for SSH access. You use a key pair to log in to your instance securely. You specify the name of the key pair when you create your compute environment, then provide the private key when you log in using SSH.

If you didn't create a key pair already, you can create one using the Amazon EC2 console. Note that, if you plan to launch instances in multiple Amazon Web Services Regions, create a key pair in each Region. For more information about Regions, see Regions and Availability Zones in the Amazon EC2 User Guide.

#### To create a key pair

- Open the Amazon EC2 console at https://console.amazonaws.cn/ec2/. 1.
- 2. From the navigation bar, select an Amazon Web Services Region for the key pair. You can select any Region that's available to you, regardless of your location: however, key pairs are specific to a Region. For example, if you plan to launch an instance in the US West (Oregon) Region, create a key pair for the instance in the same Region.
- In the navigation pane, choose **Key Pairs**, **Create Key Pair**. 3.
- In the Create Key Pair dialog box, for Key pair name, enter a name for the new key pair, and choose Create. Choose a name that you can remember, such as your user name, followed by key-pair, plus the Region name. For example, me-key-pair-uswest2.

Create a key pair

The private key file is automatically downloaded by your browser. The base file name is the 5. name that you specified as the name of your key pair, and the file name extension is .pem. Save the private key file in a safe place.

#### Important

This is the only chance for you to save the private key file. You need to provide the name of your key pair when you launch an instance and the corresponding private key each time that you connect to the instance.

If you use an SSH client on a Mac or Linux computer to connect to your Linux instance, use the following command to set the permissions of your private key file. That way, only you can read it.

```
$ chmod 400 your_user_name-key-pair-region_name.pem
```

For more information, see Amazon EC2 Key Pairs in the Amazon EC2 User Guide.

#### To connect to your instance using your key pair

To connect to your Linux instance from a computer running Mac or Linux, specify the .pem file to your SSH client with the -i option and the path to your private key. To connect to your Linux instance from a computer running Windows, use either MindTerm or PuTTY. If you plan to use PuTTY, install it and use the following procedure to convert the .pem file to a .ppk file.

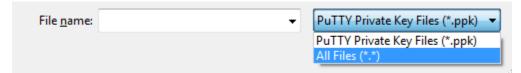
#### (Optional) To prepare to connect to a Linux instance from Windows using PuTTY

- Download and install PuTTY from http://www.chiark.greenend.org.uk/~sgtatham/putty/. Be 1. sure to install the entire suite.
- Start PuTTYgen (for example, from the Start menu, choose All Programs, PuTTY, and PuTTYgen).
- Under Type of key to generate, choose RSA. If you're using an earlier version of PuTTYgen, choose SSH-2 RSA.



Create a key pair

4. Choose **Load**. By default, PuTTYgen displays only files with the extension .ppk. To locate your .pem file, choose the option to display files of all types.



- 5. Select the private key file that you created in the previous procedure and choose **Open**. Choose **OK** to dismiss the confirmation dialog box.
- 6. Choose **Save private key**. PuTTYgen displays a warning about saving the key without a passphrase. Choose **Yes**.
- 7. Specify the same name for the key that you used for the key pair. PuTTY automatically adds the .ppk file extension.

#### Create a VPC

With Amazon Virtual Private Cloud (Amazon VPC), you can launch Amazon resources into a virtual network that you've defined. We strongly recommend that you launch your container instances in a VPC.

If you have a default VPC, you also can skip this section and move to the next task <u>Create a security group</u>. To determine whether you have a default VPC, see <u>Supported Platforms in the Amazon EC2</u> <u>Console</u> in the <u>Amazon EC2 User Guide</u>

For information about how to create an Amazon VPC, see <u>Create a VPC only</u> in the *Amazon VPC User Guide*. Refer to the following table to determine what options to select.

Option	Value
Resources to create	VPC only
Name	Optionally provide a name for your VPC.
IPv4 CIDR block	IPv4 CIDR manual input  The CIDR block size must have a size between /16 and /28.

Create a VPC 10

Option	Value
IPv6 CIDR block	No IPv6 CIDR block
Tenancy	Default

For more information about Amazon VPC, see What is Amazon VPC? in the Amazon VPC User Guide.

## Create a security group

Security groups act as a firewall for associated compute environment container instances, controlling both inbound and outbound traffic at the container instance level. A security group can be used only in the VPC for which it is created.

You can add rules to a security group that enable you to connect to your container instance from your IP address using SSH. You can also add rules that allow inbound and outbound HTTP and HTTPS access from anywhere. Add any rules to open ports that are required by your tasks.

Note that if you plan to launch container instances in multiple Regions, you need to create a security group in each Region. For more information, see Regions and Availability Zones in the Amazon EC2 User Guide.



#### Note

You need the public IP address of your local computer, which you can get using a service. For example, we provide the following service: http://checkip.amazonaws.com/ or https:// checkip.amazonaws.com/. To locate another service that provides your IP address, use the search phrase "what is my IP address." If you're connecting through an Internet service provider (ISP) or from behind a firewall without a static IP address, find out the range of IP addresses that are used by client computers.

#### To create a security group using the console

- Open the Amazon VPC console at https://console.amazonaws.cn/vpc/. 1.
- 2. In the navigation pane, choose **Security Groups**.
- Choose **Create security group**.

Create a security group

Enter a name and description for the security group. You cannot change the name and description of a security group after it is created.

- From **VPC**, choose the VPC. 5.
- (Optional) By default, new security groups start with only an outbound rule that allows all traffic to leave the resource. You must add rules to enable any inbound traffic or to restrict the outbound traffic.

Amazon Batch container instances don't require any inbound ports to be open. However, you might want to add an SSH rule. That way, you can log into the container instance and examine the containers in jobs with Docker commands. If you want your container instance to host a job that runs a web server, you can also add rules for HTTP. Complete the following steps to add these optional security group rules.

On the **Inbound** tab, create the following rules and choose **Create**:

- Choose **Add Rule**. For **Type**, choose **HTTP**. For **Source**, choose **Anywhere** (0.0.0.0/0).
- Choose Add Rule. For Type, choose SSH. For Source, choose Custom IP, and specify the public IP address of your computer or network in Classless Inter-Domain Routing (CIDR) notation. If your company allocates addresses from a range, specify the entire range, such as 203.0.113.0/24. To specify an individual IP address in CIDR notation, choose My IP. This adds the routing prefix /32 to the public IP address.



#### Note

For security reasons, we don't recommend that you allow SSH access from all IP addresses (0.0.0.0/0) to your instance but only for testing purposes and only for a short time.

- You can add tags now, or you can add them later. To add a tag, choose **Add new tag** and enter the tag key and value.
- Choose Create security group. 8.

To create a security group using the command line, see create-security-group (Amazon CLI)

For more information about security groups, see Work with security groups.

12 Create a security group

## **Install the Amazon CLI**

To use the Amazon CLI with Amazon Batch, install the latest Amazon CLI version. For information about installing the Amazon CLI or upgrading it to the latest version, see <u>Installing the Amazon</u> Command Line Interface in the Amazon Command Line Interface User Guide.

Install the Amazon CLI 13

## **Getting started with Amazon Batch tutorials**

You can use the Amazon Batch first-run wizard to get started quickly with Amazon Batch. After you complete the Prerequisites, you can use the first-run wizard to create a compute environment, a job definition, and a job queue.

You can also submit a sample "Hello World" job using the Amazon Batch first-run wizard to test your configuration. If you already have a Docker image that you want to launch in Amazon Batch, you can use that image to create a job definition.

Afterward, you can use the Amazon Batch first-run wizard to create a compute environment, job queue, and submit a sample *Hello World* job.

## Getting started with Amazon EC2 orchestration using the Wizard

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services Cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster.

You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

#### **Overview**

This tutorial demonstrates how to setup Amazon Batch with the Wizard to configure Amazon EC2 and run Hello World.

#### **Intended Audience**

This tutorial is designed for system administrators and developers responsible for setting up, testing, and deploying Amazon Batch.

#### **Features Used**

This tutorial shows you how to use the Amazon Batch console wizard to:

- Create and configure an Amazon EC2 compute environment
- Create a job queue.

- Create a job definition
- Create and submit a job to run
- View the output of the job in CloudWatch

#### **Time Required**

It should take about 10–15 minutes to complete this tutorial.

#### **Regional Restrictions**

There are no country or regional restrictions associated with using this solution.

#### **Resource Usage Costs**

There's no charge for creating an Amazon account. However, by implementing this solution, you might incur some or all of the costs that are listed in the following table.

Description	Cost (US dollars)
Amazon EC2 instance	You pay for each Amazon EC2 instance that is created. For more information about pricing, see <a href="Mazon EC2 Pricing"><u>Amazon EC2 Pricing</u></a> .

## **Prerequisites**

Before you begin:

- Create an Amazon Web Services account if you don't have one.
- Create the ecsInstanceRole Instance role.

## **Step 1: Create a compute environment**



#### Important

To get started as simply and quickly as possible, this tutorial includes steps with default settings. Before creating for production use, we recommend that you familiarize yourself with all settings and deploy with the settings that meet your requirements.

To create a compute environment for an Amazon EC2 orchestration, do the following:

Prerequisites 15

- Open the Amazon Batch console first-run wizard.
- 2. For Configure job and orchestration type, choose Amazon Elastic Compute Cloud(Amazon EC2).
- Choose Next.
- 4. In the **Compute environment configuration** section for **Name**, specify a unique name for your compute environment. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 5. For **Instance role**, choose an existing instance role that has the required IAM permissions attached. This instance role allows the Amazon ECS container instances in your compute environment to make calls to the required Amazon API operations. For more information, see Amazon ECS instance role.

The default name of the **Instance role** is ecsInstanceRole.

- 6. For **Instance configuration** you can leave the default settings.
- 7. For **Network configuration** use your default VPC for the Amazon Web Services Region.
- 8. Choose Next.

## **Step 2: Create a job queue**

A job queue stores your submitted jobs until the Amazon Batch Scheduler runs the job on a resource in your compute environment. For more information, see *Job queues* 

To create a job queue for an Amazon EC2 orchestration, do the following:

- 1. For **Job queue configuration** for **Name**, specify a unique name for your job queue. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 2. For all other configuration options you can leave the default value.
- 3. Choose **Next**.

## **Step 3: Create a job definition**

Amazon Batch job definitions specify how jobs are to be run. Even though each job must reference a job definition, many of the parameters that are specified in the job definition can be overridden at runtime.

Step 2: Create a job queue 16

To create the job definition:

#### 1. For Create a job definition

for Name, specify a unique name for your job queue. The name can be up to 128
characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-),
and underscores ( ).

- b. For **Command optional** you can change hello world to a custom message or leave it as is.
- 2. For all other configuration options you can leave the default value.
- Choose Next.

## Step 4: Create a job

To create a job, do the following:

- 1. In the **Job configuration** section for **Name**, specify a unique name for the job. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 2. For all other configuration options you can leave the default value.
- 3. Choose Next.

## **Step 5: Review and create**

On the **Review and create** page, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create resources**.

- For Review and create choose Create resources.
- 2. A window opens as Amazon Batch starts to allocate your resources. Once complete choose **Go to dashboard**. On the dashboard you should see all of your allocated resources and that the job is in the Runnable state. Your job is scheduled to run and should complete in 2–3 minuets.

## **Step 6: View the Job's output**

To view the Job's output, do the following:

Step 4: Create a job

- 1. In the navigation pane choose **Jobs**.
- 2. In the Job queue drop down choose the Job queue you created for the tutorial.
- 3. The **Jobs** table lists all of your Jobs and what their current status is. Once the Job's **Status** is **Succeeded** choose the **Name** of the Job to view the Job's details.

4. In the **Details** pane choose **Log stream name**. The CloudWatch console for the Job will open and there should be one event with the **Message** of hello world or your custom message.

## Step 7: Clean up your tutorial resources

You are charged for the Amazon EC2 instance while it is enabled. You can delete the instance to stop incurring charges.

To delete the resources you created, do the following:

- 1. In the navigation pane choose **Job queue**.
- 2. In the **Job queue** table choose the Job queue you created for the tutorial.
- 3. Choose **Disable**. Once the Job queue **State** is Disabled you can choose **Delete**.
- 4. Once the Job queue is deleted, in the navigation pane choose **Compute environments**.
- 5. Choose the compute environment you created for this tutorial and then choose **Disable**. It may take 1–2 minuets for the compute environment to complete being disabled.
- Once the compute environment's State is Disabled, choose Delete. It may take 1–2 minuets for the compute environment to be deleted.

### **Additional resources**

After you complete the tutorial, you might want to explore the following topics::

- Explore the Amazon Batch core components. For more information, see <u>Components of Amazon</u> Batch.
- Learn more about the different Compute Environments available in Amazon Batch.
- Learn more about <u>Job queues</u> and their different scheduling options.
- Learn more about <u>Job definitions</u> and the different configuration options.
- Learn more about the different types of <u>Jobs</u>.

# Getting started with Amazon Batch and Fargate orchestration using the Wizard

Amazon Fargate launches and scales the compute to closely match the resource requirements that you specify for the container. With Fargate, you don't need to over-provision or pay for additional servers. For more information, see Fargate.

#### **Overview**

This tutorial demonstrates how to setup Amazon Batch with the Wizard to configure Amazon Fargate and run Hello World.

#### **Intended Audience**

This tutorial is designed for system administrators and developers responsible for setting up, testing, and deploying Amazon Batch.

#### **Features Used**

This tutorial shows you how to use the Amazon Batch console wizard to:

- · Create and configure an Amazon Fargate compute environment
- Create a job queue.
- Create a job definition
- Create and submit a job to run
- View the output of the job in CloudWatch

#### **Time Required**

It should take about 10–15 minutes to complete this tutorial.

#### **Regional Restrictions**

There are no country or regional restrictions associated with using this solution.

#### **Resource Usage Costs**

There's no charge for creating an Amazon account. However, by implementing this solution, you might incur some or all of the costs that are listed in the following table.

Description	Cost (US dollars)
Pricing is based on requested vCPU, memory, Operating Systems, CPU Architecture, and storage resources for the Task or Pod.	For more information about pricing, see <a href="Fargate Pricing">Fargate Pricing</a> .

## **Prerequisites**

Before you begin:

- Create an Amazon Web Services account if you don't have one.
- Create the task execution role. If you haven't already created the Task Execution Role then you can create it as part of this tutorial.

## Step 1: Create a compute environment



#### Important

To get started as simply and quickly as possible, this tutorial includes steps with default settings. Before creating for production use, we recommend that you familiarize yourself with all settings and deploy with the settings that meet your requirements.

To create a compute environment for a Fargate orchestration, do the following:

- 1. Open the Amazon Batch console first-run wizard.
- 2. For **Configure job and orchestration type**, choose **Fargate**.
- 3. Choose Next.
- In the **Compute environment configuration** section for **Name**, specify a unique name for your compute environment. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 5. For all other configuration options you can leave the default value.
- Choose Next.

Prerequisites 20

## Step 2: Create a job queue

A job queue stores your submitted jobs until the Amazon Batch Scheduler runs the job on a resource in your compute environment. To create a job queue:

To create a job queue for a Fargate orchestration, do the following:

- 1. In the **Job queue configuration** section for **Name**, specify a unique name for your compute environment. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 2. For **Priority**, enter 900 for the job queue.
- 3. For all other configuration options you can leave the default value.
- 4. Choose Next.

## **Step 3: Create a job definition**

To create the job definition:

- 1. In the **General configuration** section:
  - In the General configuration section for Name, specify a unique name for your compute environment. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 2. In the **Fargate platform configuration** section:
  - a. Turn on **Assign public IP** to assign a public IP address. You need a public IP to download the container image unless you've setup a private image repository.
  - b. For Execution role, choose a task execution role that lets Amazon Elastic Container Service (Amazon ECS) agents make Amazon calls on your behalf. Choose either ecsTaskExecutionRole or BatchEcsTaskExecutionRole.

To create the **Execution role** choose **Create an execution role**. In the **Create IAM role** modal choose **Create IAM role**.

- i. The IAM console has the permission setting already configured for creating the execution role.
- ii. For **Trusted entity type** verify that **Amazon service** is selected.

Step 2: Create a job queue 21

- iii. For Service or user case verify that Elastic Container Service is selected.
- iv. Choose **Next**.
- v. For **Permissions policies** verify that **AmazonECSTaskExecutionRolePolicy** is selected.
- vi. Choose **Next**.
- vii. For **Name, review, and create** verify that the role name is **BatchEcsTaskExecutionRole**.
- viii. Choose Create role.
- ix. In the Amazon Batch console choose the refresh button next to **Execution role**. Choose the **BatchEcsTaskExecutionRole** execution role.
- 3. In the **Container configuration** section:
  - For Command, you can change hello world to a custom message or leave it as is.
- 4. For all other configuration options you can leave the default value.
- Choose Next.

## **Step 4: Create a job**

To create a Fargate job, do the following:

- 1. In the **Job configuration** section for **Name**, specify a unique name for the job. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 2. For all other configuration options you can leave the default value.
- Choose Next.

## **Step 5: Review and create**

On the **Review and create** page, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create resources**.

## Step 6: View the Job's output

To view the Job's output, do the following:

1. In the navigation pane choose **Jobs**.

Step 4: Create a job 22

- 2. In the **Job gueue** drop down choose the Job gueue you created for the tutorial.
- 3. The **Jobs** table lists all of your Jobs and what their current status is. Once the Job's **Status** is **Succeeded** choose the **Name** of the Job to view the Job's details.

4. In the **Details** pane choose **Log stream name**. The CloudWatch console for the Job will open and there should be one event with the **Message** of hello world or your custom message.

## Step 7: Clean up your tutorial resources

You are charged for the Amazon EC2 instance while it is enabled. You can delete the instance to stop incurring charges.

To delete the resources you created, do the following:

- 1. In the navigation pane choose **Job queue**.
- 2. In the Job queue table choose the Job queue you created for the tutorial.
- 3. Choose **Disable**. Once the Job queue **State** is Disabled you can choose **Delete**.
- 4. Once the Job queue is deleted, in the navigation pane choose **Compute environments**.
- 5. Choose the compute environment you created for this tutorial and then choose **Disable**. It may take 1–2 minuets for the compute environment to complete being disabled.
- Once the compute environment's State is Disabled, choose Delete. It may take 1–2 minuets for the compute environment to be deleted.

#### **Additional resources**

After you complete the tutorial, you might want to explore the following topics::

- Learn more about the <u>Best practices</u>.
- Explore the Amazon Batch core components. For more information, see <u>Components of Amazon</u> Batch.
- Learn more about the different Compute Environments available in Amazon Batch.
- Learn more about Job queues and their different scheduling options.
- Learn more about Job definitions and the different configuration options.
- Learn more about the different types of <u>Jobs</u>.

## Getting started with Amazon Batch and Fargate using the Amazon CLI

This tutorial demonstrates how to set up Amazon Batch with Amazon Fargate orchestration and run a simple "Hello World" job using the Amazon Command Line Interface (Amazon CLI). You'll learn how to create compute environments, job queues, job definitions, and submit jobs to Amazon Batch.

#### **Topics**

- Prerequisites
- Create an IAM execution role
- Create a compute environment
- Create a job queue
- Create a job definition
- Submit and monitor a job
- View job output
- Clean up resources
- Going to production
- Next steps

## **Prerequisites**

Before you begin this tutorial, make sure you have the following.

- 1. The Amazon CLI. If you need to install it, follow the <u>Amazon CLI installation guide</u>. You can also <u>use Amazon CloudShell</u>, which includes the Amazon CLI.
- 2. Configured your Amazon CLI with appropriate credentials. Run aws configure if you haven't set up your credentials yet.
- 3. Basic familiarity with command line interfaces and containerization concepts.
- 4. <u>How Amazon Batch works with IAM</u> to create and manage Amazon Batch resources, IAM roles, and VPC resources in your Amazon Web Services account.
- 5. A subnet ID and security group ID from a VPC in your Amazon Web Services account. If you don't have a VPC, you can create one. For more information about using the Amazon CLI to

retrieve these resource IDs, see <u>describe-subnets</u> and <u>describe-security-groups</u> in the *Amazon CLI Command Reference*.

**Time Required**: Approximately 15-20 minutes to complete this tutorial.

**Cost**: This tutorial uses Fargate compute resources. The estimated cost for completing this tutorial is less than \$0.01 USD, assuming you follow the cleanup instructions to delete resources immediately after completion. Fargate pricing is based on vCPU and memory resources consumed, charged per second with a 1-minute minimum. For current pricing information, see <a href="Amazon"><u>Amazon</u></a> Fargate pricing.

#### Create an IAM execution role

Amazon Batch requires an execution role that allows Amazon Elastic Container Service (Amazon ECS) agents to make Amazon API calls on your behalf. This role is necessary for Fargate tasks to pull container images and write logs to Amazon CloudWatch.

#### Create a trust policy document

First, create a trust policy that allows the Amazon ECS tasks service to assume the role.

#### Create the execution role

The following command creates an IAM role named BatchEcsTaskExecutionRoleTutorial using the trust policy you just created.

Create an IAM execution role 25

```
aws iam create-role \
    --role-name BatchEcsTaskExecutionRoleTutorial \
    --assume-role-policy-document file://batch-execution-role-trust-policy.json
```

The following shows how the output looks when the command runs successfully.

```
{
    "Role": {
        "Path": "/",
        "RoleName": "BatchEcsTaskExecutionRoleTutorial",
        "RoleId": "AROAUVBF026T7xmpl3RGYTV0",
        "Arn": "arn:aws:iam::123456789012:role/BatchEcsTaskExecutionRoleTutorial",
        "CreateDate": "2025-01-13T17:34:28+00:00",
        "AssumeRolePolicyDocument": {
            "Version": "2012-10-17",
            "Statement": [
                {
                    "Effect": "Allow",
                    "Principal": {
                         "Service": "ecs-tasks.amazonaws.com"
                    },
                    "Action": "sts:AssumeRole"
                }
            ]
        }
    }
}
```

### Attach the required policy

Attach the Amazon managed policy that provides the necessary permissions for Amazon ECS task execution.

```
aws iam attach-role-policy \
    --role-name BatchEcsTaskExecutionRoleTutorial \
    --policy-arn arn:aws:iam::aws:policy/service-role/AmazonECSTaskExecutionRolePolicy
```

The role is now ready to be used by Amazon Batch for Fargate task execution.

Create an IAM execution role 26

# Create a compute environment

A compute environment defines the compute resources where your batch jobs will run. For this tutorial, you'll create a managed Fargate compute environment that automatically provisions and scales resources based on job requirements.

#### Create the compute environment

The following command creates a Fargate compute environment. Replace the example subnet and security group IDs with your own per the Prerequisites.

```
aws batch create-compute-environment \
    --compute-environment-name my-fargate-compute-env \
    --type MANAGED \
    --state ENABLED \
    --compute-resources type=FARGATE,maxvCpus=128,subnets=subnet-
a123456b,securityGroupIds=sg-a12b3456
```

The following shows how the output looks when the command runs successfully.

```
{
    "computeEnvironmentName": "my-fargate-compute-env",
    "computeEnvironmentArn": "arn:aws:batch:us-west-2:123456789012:compute-environment/
my-fargate-compute-env"
}
```

#### Wait for the compute environment to be ready

Check the status of your compute environment to ensure it's ready before proceeding.

```
aws batch describe-compute-environments \
    --compute-environments my-fargate-compute-env \
    --query 'computeEnvironments[0].status'
```

```
"VALID"
```

When the status shows VALID, your compute environment is ready to accept jobs.

# Create a job queue

A job queue stores submitted jobs until the Amazon Batch scheduler runs them on resources in your compute environment. Jobs are processed in priority order within the queue.

# Create the job queue

The following command creates a job queue with priority 900 that uses your Fargate compute environment.

```
aws batch create-job-queue \
    --job-queue-name my-fargate-job-queue \
    --state ENABLED \
    --priority 900 \
    --compute-environment-order order=1,computeEnvironment=my-fargate-compute-env
```

The following shows how the output looks when the command runs successfully.

```
{
    "jobQueueName": "my-fargate-job-queue",
    "jobQueueArn": "arn:aws:batch:us-west-2:123456789012:job-queue/my-fargate-job-queue"
}
```

### Verify the job queue is ready

Check that your job queue is in the ENABLED state and ready to accept jobs.

```
aws batch describe-job-queues \
    --job-queues my-fargate-job-queue \
    --query 'jobQueues[0].state' "ENABLED"
```

# Create a job definition

A job definition specifies how jobs are to be run, including the Docker image to use, resource requirements, and other parameters. For Fargate, you'll use resource requirements instead of traditional vCPU and memory parameters.

# Create the job definition

Create a job queue 28

The following command creates a job definition that runs a simple "hello world" command using the busybox container image. Replace 123456789012 with your actual Amazon Web Services account ID and replace the example Amazon Web Services Region with your own.

```
aws batch register-job-definition \
    --job-definition-name my-fargate-job-def \
    --type container \
    --platform-capabilities FARGATE \
    --container-properties '{
        "image": "busybox",
        "resourceRequirements": [
            {"type": "VCPU", "value": "0.25"},
            {"type": "MEMORY", "value": "512"}
        ],
        "command": ["echo", "hello world"],
        "networkConfiguration": {
            "assignPublicIp": "ENABLED"
        },
        "executionRoleArn": "arn:aws:iam::123456789012:role/
BatchEcsTaskExecutionRoleTutorial"
    },
{
    "jobDefinitionName": "my-fargate-job-def",
    "jobDefinitionArn": "arn:aws:batch:us-west-2:123456789012:job-definition/my-
fargate-job-def:1",
    "revision": 1
}'
```

The job definition specifies 0.25 vCPU and 512 MB of memory, which are the minimum resources for a Fargate task. The assignPublicIp setting is enabled so the container can pull the busybox image from Docker Hub.

# Submit and monitor a job

Now that you have all the necessary components, you can submit a job to your queue and monitor its progress.

## Submit a job

The following command submits a job to your queue using the job definition you created.

```
aws batch submit-job \
```

Submit and monitor a job 29

```
--job-name my-hello-world-job \
--job-queue my-fargate-job-queue \
--job-definition my-fargate-job-def
```

The following shows how the output looks when the command runs successfully.

```
{
    "jobArn": "arn:aws:batch:us-west-2:123456789012:job/my-hello-world-job",
    "jobName": "my-hello-world-job",
    "jobId": "1509xmpl-4224-4da6-9ba9-1d1acc96431a"
}
```

Make note of the jobId returned in the response, as you'll use it to monitor the job's progress.

#### **Monitor job status**

Use the job ID to check the status of your job. The job will progress through several states: SUBMITTED, PENDING, RUNNABLE, STARTING, RUNNING, and finally SUCCEEDED or FAILED.

```
aws batch describe-jobs --jobs 1509xmpl-4224-4da6-9ba9-1d1acc96431a
```

The following shows how the output looks when the command runs successfully.

When the status shows SUCCEEDED, your job has completed successfully.

Submit and monitor a job 30

# View job output

After your job completes, you can view its output in Amazon CloudWatch Logs.

#### Get the log stream name

First, retrieve the log stream name from the job details. Replace the example job ID with your own.

```
aws batch describe-jobs --jobs 1509xmpl-4224-4da6-9ba9-1d1acc96431a \
    --query 'jobs[0].attempts[0].containers[0].logStreamName' \
    --output text

my-fargate-job-def/default/1509xmpl-4224-4da6-9ba9-1d1acc96431a
```

#### View the job logs

Use the log stream name to retrieve the job's output from CloudWatch Logs.

```
aws logs get-log-events \
    --log-group-name /aws/batch/job \
    --log-stream-name my-fargate-job-def/default/1509xmpl-4224-4da6-9ba9-1d1acc96431a \
    --query 'events[*].message' \
    --output text
```

The output shows "hello world", confirming that your job ran successfully.

# Clean up resources

To avoid ongoing charges, clean up the resources you created in this tutorial. You must delete resources in the correct order due to dependencies.

### Disable and delete the job queue

First, disable the job queue, then delete it.

```
aws batch update-job-queue \
--job-queue my-fargate-job-queue \
--state DISABLED

aws batch delete-job-queue \
--job-queue my-fargate-job-queue
```

View job output 31

#### Disable and delete the compute environment

After the job queue is deleted, disable and delete the compute environment.

```
aws batch update-compute-environment \
    --compute-environment my-fargate-compute-env \
    --state DISABLED

aws batch delete-compute-environment \
    --compute-environment my-fargate-compute-env
```

#### Clean up the IAM role

Remove the policy attachment and delete the IAM role.

```
aws iam detach-role-policy \
    --role-name BatchEcsTaskExecutionRoleTutorial \
    --policy-arn arn:aws:iam::aws:policy/service-role/AmazonECSTaskExecutionRolePolicy

aws iam delete-role \
    --role-name BatchEcsTaskExecutionRoleTutorial
```

#### Remove temporary files

Delete the trust policy file you created.

```
rm batch-execution-role-trust-policy.json
```

All resources have been successfully cleaned up.

# **Going to production**

This tutorial is designed to help you understand how Amazon Batch works with Fargate. For production deployments, consider the following additional requirements:

## **Security considerations:**

 Create dedicated security groups with minimal required access instead of using default security groups

Going to production 32

• Use private subnets with NAT Gateway instead of public IP assignment for containers

- · Store container images in Amazon ECR instead of using public repositories
- Implement VPC endpoints for Amazon service communication to avoid internet traffic

#### **Architecture considerations:**

- Deploy across multiple Availability Zones for high availability
- Implement job retry strategies and dead letter queues for error handling
- Use multiple job queues with different priorities for workload management
- Configure auto scaling policies based on queue depth and resource utilization
- Implement monitoring and alerting for job failures and resource utilization

#### **Operational considerations:**

- Set up CloudWatch dashboards and alarms for monitoring
- Implement proper logging and audit trails
- Use Amazon CloudFormation or the Amazon CDK for infrastructure as code
- Establish backup and disaster recovery procedures

For comprehensive guidance on production-ready architectures, see the <u>Amazon Well-Architected</u> <u>Framework and Amazon Security Best Practices</u>.

# **Next steps**

Now that you've completed this tutorial, you can explore more advanced Amazon Batch features:

- Job queues Learn about job queue scheduling and priority management
- <u>Job definitions</u> Explore advanced job definition configurations including environment variables, volumes, and retry strategies
- <u>Compute environments for Amazon Batch</u> Understand different compute environment types and scaling options
- Multi-node parallel jobs Run jobs that span multiple compute nodes
- Array jobs Submit large numbers of similar jobs efficiently

Next steps 33

• Best practices for Amazon Batch – Learn optimization techniques for production workloads

# **Getting started with Amazon Batch on Amazon EKS**

Amazon Batch on Amazon EKS is a managed service for scheduling and scaling batch workloads into existing Amazon EKS clusters. Amazon Batch doesn't create, administer, or perform lifecycle operations of your Amazon EKS clusters on your behalf. Amazon Batch orchestration scales up and down nodes managed by Amazon Batch and run pods on those nodes.

Amazon Batch doesn't touch nodes, auto scaling node groups or pods lifecycles that aren't associated with Amazon Batch compute environments within your Amazon EKS cluster. For Amazon Batch to operate effectively, its <u>service-linked role</u> needs Kubernetes role-based access control (RBAC) permissions in your existing Amazon EKS cluster. For more information, see <u>Using RBAC Authorization</u> in the *Kubernetes documentation*.

Amazon Batch requires a Kubernetes namespace where it can scope pods as Amazon Batch jobs into. We recommend a dedicated namespace to isolate the Amazon Batch pods from your other cluster workloads.

After Amazon Batch has been given RBAC access and a namespace has been established, you can associate that Amazon EKS cluster to an Amazon Batch compute environment using the <a href="CreateComputeEnvironment">CreateComputeEnvironment</a> API operation. A job queue can be associated with this new Amazon EKS compute environment. Amazon Batch jobs are submitted to the job queue based on an Amazon EKS job definition using the <a href="SubmitJob">SubmitJob</a> API operation. Amazon Batch then launches Amazon Batch managed nodes and place jobs from job queue as Kubernetes pods into the EKS cluster associated with an Amazon Batch compute environment.

The following sections cover how to get set up for Amazon Batch on Amazon EKS.

#### **Contents**

- Overview
- Prerequisites
- Step 1: Create your Amazon EKS cluster for Amazon Batch
- Step 2: Prepare your Amazon EKS cluster for Amazon Batch
- Step 3: Create an Amazon EKS compute environment
- Step 4: Create a job queue and attach the compute environment

- Step 5: Create a job definition
- Step 6: Submit a job
- Step 7: View the Job's output
- Step 8: (Optional) Submit a job with overrides
- Step 9: Clean up your tutorial resources
- Additional resources

#### **Overview**

This tutorial demonstrates how to setup Amazon Batch with Amazon EKS using the Amazon CLI, kubectl and eksctl.

#### **Intended Audience**

This tutorial is designed for system administrators and developers responsible for setting up, testing, and deploying Amazon Batch.

#### **Features Used**

This tutorial shows you how to use the Amazon CLI, to:

- Create and configure an Amazon EKS compute environment
- Create a job queue.
- Create a job definition
- Create and submit a job to run
- Submit a job with overrides

# **Time Required**

It should take about 30–40 minutes to complete this tutorial.

# **Regional Restrictions**

There are no country or regional restrictions associated with using this solution.

# **Resource Usage Costs**

There's no charge for creating an Amazon account. However, by implementing this solution, you might incur some or all of the costs that are listed in the following table.

Overview 35

Description	Cost (US dollars)
You are charged by the cluster hour	Varies depending on Instance, see <u>Amazon</u> <u>EKS pricing</u>

# **Prerequisites**

Before starting this tutorial, you must install and configure the following tools and resources that you need to create and manage both Amazon Batch and Amazon EKS resources.

- Amazon CLI A command line tool for working with Amazon services, including Amazon EKS.
   This guide requires that you use version 2.8.6 or later or 1.26.0 or later. For more information, see <u>Installing</u>, <u>updating</u>, <u>and uninstalling the Amazon CLI</u> in the *Amazon Command Line Interface User Guide*. After installing the Amazon CLI, we recommend that you also configure it. For more information, see <u>Quick configuration with aws configure</u> in the *Amazon Command Line Interface User Guide*.
- **kubect1** A command line tool for working with Kubernetes clusters. This guide requires that you use version 1.23 or later. For more information, see <u>Installing or updating kubect1</u> in the *Amazon EKS User Guide*.
- eksct1 A command line tool for working with Amazon EKS clusters that automates many
  individual tasks. This guide requires that you use version 0.115.0 or later. For more information,
  see Installing or updating eksct1 in the Amazon EKS User Guide.
- Required IAM permissions The IAM security principal that you're using must have permissions
  to work with Amazon EKS IAM roles and service linked roles, Amazon CloudFormation, and a
  VPC and related resources. For more information, see <u>Actions, resources, and condition keys for Amazon Elastic Kubernetes Service</u> and <u>Using service-linked roles</u> in the *IAM User Guide*. You
  must complete all steps in this guide as the same user.
- Permissions Users calling the <u>CreateComputeEnvironment</u> API operation to create
  a compute environment that uses Amazon EKS resources require permissions to the
  eks:DescribeCluster API operation.
- Amazon Web Services account number You need to know your Amazon Web Services account ID. Follow the directions in Viewing your Amazon Web Services account ID.
- (Optional) CloudWatch To examine the details of (Optional) Submit a job with overrides, logging must be configured. For more information, see Use CloudWatch Logs to monitor Amazon Batch on Amazon EKS jobs.

Prerequisites 36

# Step 1: Create your Amazon EKS cluster for Amazon Batch

#### Important

To get started as simply and quickly as possible, this tutorial includes steps with default settings. Before creating for production use, we recommend that you familiarize yourself with all settings and deploy with the settings that meet your requirements.

Once you have installed the prerequisites you need to create your cluster using eksctl. Creating the cluster can take between 10-15 minutes.

```
eksctl create cluster --name my-cluster-name --region region-code
```

In the preceding command replace:

- Replace my-cluster-name with the name you want to use for your cluster.
- Replace <u>region-code</u> with the Amazon Web Services Region to create the cluster in, for example us-west-2.

The cluster name and region are needed for later in this tutorial.

# Step 2: Prepare your Amazon EKS cluster for Amazon Batch

All steps are required.

# Create a dedicated namespace for Amazon Batch jobs

Use kubect1 to create a new namespace.

```
$ namespace=my-aws-batch-namespace
```

```
$ cat - <<EOF | kubectl create -f -</pre>
  "apiVersion": "v1",
  "kind": "Namespace",
  "metadata": {
    "name": "${namespace}",
```

```
"labels": {
        "name": "${namespace}"
     }
}

EOF
```

#### Output:

```
namespace/my-aws-batch-namespace created
```

#### 2. Enable access via role-based access control (RBAC)

Use kubectl to create a Kubernetes role for the cluster to allow Amazon Batch to watch nodes and pods, and to bind the role. You must do this once for each EKS cluster.

```
$ cat - <<EOF | kubectl apply -f -</pre>
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: aws-batch-cluster-role
rules:
  - apiGroups: [""]
    resources: ["namespaces"]
    verbs: ["get"]
  - apiGroups: [""]
    resources: ["nodes"]
    verbs: ["get", "list", "watch"]
  - apiGroups: [""]
    resources: ["pods"]
    verbs: ["get", "list", "watch"]
  - apiGroups: [""]
    resources: ["events"]
    verbs: ["list"]
  - apiGroups: [""]
    resources: ["configmaps"]
    verbs: ["get", "list", "watch"]
  - apiGroups: ["apps"]
    resources: ["daemonsets", "deployments", "statefulsets", "replicasets"]
    verbs: ["get", "list", "watch"]
  - apiGroups: ["rbac.authorization.k8s.io"]
    resources: ["clusterroles", "clusterrolebindings"]
    verbs: ["get", "list"]
```

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
    name: aws-batch-cluster-role-binding
subjects:
    - kind: User
    name: aws-batch
    apiGroup: rbac.authorization.k8s.io
roleRef:
    kind: ClusterRole
    name: aws-batch-cluster-role
apiGroup: rbac.authorization.k8s.io
EOF
```

#### **Output:**

```
clusterrole.rbac.authorization.k8s.io/aws-batch-cluster-role created clusterrolebinding.rbac.authorization.k8s.io/aws-batch-cluster-role-binding created
```

3. Create namespace-scoped Kubernetes role for Amazon Batch to manage and lifecycle pods and bind it. You must do this once for each unique namespace.

```
$ namespace=my-aws-batch-namespace
```

```
$ cat - <<EOF | kubectl apply -f - --namespace "${namespace}"</pre>
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: aws-batch-compute-environment-role
  namespace: ${namespace}
rules:
  - apiGroups: [""]
    resources: ["pods"]
    verbs: ["create", "get", "list", "watch", "delete", "patch"]
  - apiGroups: [""]
    resources: ["serviceaccounts"]
    verbs: ["get", "list"]
  - apiGroups: ["rbac.authorization.k8s.io"]
    resources: ["roles", "rolebindings"]
    verbs: ["get", "list"]
```

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
    name: aws-batch-compute-environment-role-binding
    namespace: ${namespace}
subjects:
    - kind: User
    name: aws-batch
    apiGroup: rbac.authorization.k8s.io
roleRef:
    kind: Role
    name: aws-batch-compute-environment-role
    apiGroup: rbac.authorization.k8s.io
EOF
```

#### Output:

```
role.rbac.authorization.k8s.io/aws-batch-compute-environment-role created rolebinding.rbac.authorization.k8s.io/aws-batch-compute-environment-role-binding created
```

4. Update Kubernetes aws-auth configuration map to map the preceding RBAC permissions to the Amazon Batch service-linked role.

In the following command replace:

Replace < your-account-number > with your Amazon Web Services account number.

```
$ eksctl create iamidentitymapping \
    --cluster my-cluster-name \
    --arn "arn:aws-cn:iam::<your-account-number>:role/AWSServiceRoleForBatch" \
    --username aws-batch
```

#### Output:

```
2022-10-25 20:19:57 [#] adding identity "arn:aws-cn:iam::
number>:role/AWSServiceRoleForBatch" to auth ConfigMap
```



#### Note

The path aws-service-role/batch.amazonaws.com/ has been removed from the ARN of the service-linked role. This is because of an issue with the aws-auth configuration map. For more information, see Roles with paths don't work when the path is included in their ARN in the aws-authconfigmap.

# Step 3: Create an Amazon EKS compute environment

Amazon Batch compute environments define compute resource parameters to meet your batch workload needs. In a managed compute environment, Amazon Batch helps you to manage the capacity and instance types of the compute resources (Kubernetes nodes) within your Amazon EKS cluster. This is based on the compute resource specification that you define when you create the compute environment. You can use EC2 On-Demand Instances or EC2 Spot Instances.

Now that the AWSServiceRoleForBatch service-linked role has access to your Amazon EKS cluster, you can create Amazon Batch resources. First, create a compute environment that points to your Amazon EKS cluster.

- For subnets run eksctl get cluster my-cluster-name to get the subnets used by the cluster.
- For securityGroupIds parameter you can use the same security group as the Amazon EKS cluster. This command retrieves the security group ID for the cluster.

```
$ aws eks describe-cluster \
    -- name my-cluster-name \
    --query cluster.resourcesVpcConfig.clusterSecurityGroupId
```

 The instanceRole is created when you create the cluster. To find the instanceRole list all entities that use the AmazonEKSWorkerNodePolicy policy:

```
aws iam list-entities-for-policy --policy-arn arn:aws:iam::aws:policy/
AmazonEKSWorkerNodePolicy
```

The name of the policy role contains the name of the cluster that you created eksctl-my*cluster-name*-nodegroup-example.

To find the instanceRole arn run the following command:

```
$ aws iam list-instance-profiles-for-role --role-name eksctl-my-cluster-name-
nodegroup-example
```

#### Output:

```
INSTANCEPROFILES arn:aws:iam::<your-account-number>:instance-profile/
eks-04cb2200-94b9-c297-8dbe-87f12example
```

For more information, see <u>Creating the Amazon EKS node IAM role</u> and <u>Enabling IAM principal access to your cluster</u> in the <u>Amazon EKS User Guide</u>. If you're using pod networking, see <u>Configuring the Amazon VPC CNI plugin for Kubernetes to use IAM roles for service accounts</u> in the Amazon EKS User Guide.

```
$ cat <<EOF > ./batch-eks-compute-environment.json
{
  "computeEnvironmentName": "My-Eks-CE1",
  "type": "MANAGED",
  "state": "ENABLED",
  "eksConfiguration": {
    "eksClusterArn": "arn:aws-cn:eks:region-code:your-account-number:cluster/my-
cluster-name",
    "kubernetesNamespace": "my-aws-batch-namespace"
  },
  "computeResources": {
    "type": "EC2",
    "allocationStrategy": "BEST_FIT_PROGRESSIVE",
    "minvCpus": 0,
    "maxvCpus": 128,
    "instanceTypes": [
        "m5"
    ],
    "subnets": [
        "<eks-cluster-subnets-with-access-to-internet-for-image-pull>"
    ],
    "securityGroupIds": [
        "<eks-cluster-sq>"
    "instanceRole": "<eks-instance-profile>"
```

```
}
}
E<sub>0</sub>F
```

```
$ aws batch create-compute-environment --cli-input-json file://./batch-eks-compute-
environment.json
```

#### **Notes**

 Maintenance of an Amazon EKS compute environment is a shared responsibility. For more information, see Shared responsibility of the Kubernetes nodes.

# Step 4: Create a job queue and attach the compute environment

#### Important

It's important to confirm that the compute environment is healthy before proceeding. The DescribeComputeEnvironments API operation can be used to do this.

```
$ aws batch describe-compute-environments --compute-environments My-Eks-CE1
```

Confirm that the status parameter is not INVALID. If it is, look at the statusReason parameter for the cause. For more information, see Troubleshooting Amazon Batch.

Jobs submitted to this new job queue are run as pods on Amazon Batch managed nodes that joined the Amazon EKS cluster that's associated with your compute environment.

```
$ cat <<EOF > ./batch-eks-job-queue.json
 {
    "jobQueueName": "My-Eks-JQ1",
    "priority": 10,
    "computeEnvironmentOrder": [
        "order": 1,
        "computeEnvironment": "My-Eks-CE1"
      }
    ]
  }
```

**EOF** 

```
$ aws batch create-job-queue --cli-input-json file://./batch-eks-job-queue.json
```

# **Step 5: Create a job definition**

The following Job definition instructs the pod to sleep for 60 seconds.

```
$ cat <<EOF > ./batch-eks-job-definition.json
  "jobDefinitionName": "MyJobOnEks_Sleep",
  "type": "container",
  "eksProperties": {
    "podProperties": {
      "hostNetwork": true,
      "containers": [
          "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
          "command": [
            "sleep",
            "60"
          ],
          "resources": {
            "limits": {
              "cpu": "1",
              "memory": "1024Mi"
            }
          }
        }
      ],
      "metadata": {
        "labels": {
          "environment": "test"
        }
      }
    }
  }
}
E0F
```

```
$ aws batch register-job-definition --cli-input-json file://./batch-eks-job-
definition.json
```

#### **Notes**

There are considerations for the cpu and memory parameters. For more information, see
 Memory and vCPU considerations for Amazon Batch on Amazon EKS.

# Step 6: Submit a job

Run the following Amazon CLI command to submit a new Job.

```
$ aws batch submit-job --job-queue My-Eks-JQ1 \
    --job-definition MyJobOnEks_Sleep --job-name My-Eks-Job1
```

To check the status of a Job:

```
$ aws batch describe-jobs --job <jobId-from-submit-response>
```

#### **Notes**

For more information about running jobs on Amazon EKS resources, see <u>Amazon EKS jobs</u>.

# Step 7: View the Job's output

To view the Job's output, do the following:

- 1. Open the Amazon Batch console at <a href="https://console.amazonaws.cn/batch/">https://console.amazonaws.cn/batch/</a>.
- 2. In the navigation pane choose **Jobs**.
- 3. In the **Job queue** drop down choose the Job queue you created for the tutorial.
- 4. The **Jobs** table lists all of your Jobs and what their current status is. Once the Job's **Status** is **Succeeded** choose the **Name** of the Job, *My-Eks-JQ1*, to view the Job's details.
- 5. In the **Details** pane the **Started at** and **Stopped at** times should be one minute apart.

Step 6: Submit a job 45

# Step 8: (Optional) Submit a job with overrides

This job overrides the command passed to the container. Amazon Batch aggressively cleans up the pods after the jobs complete to reduce the load to Kubernetes. To examine the details of a job, logging must be configured. For more information, see <u>Use CloudWatch Logs to monitor Amazon</u> Batch on Amazon EKS jobs.

```
$ cat <<EOF > ./submit-job-override.json
{
  "jobName": "EksWithOverrides",
  "jobQueue": "My-Eks-JQ1",
  "jobDefinition": "MyJobOnEks_Sleep",
  "eksPropertiesOverride": {
    "podProperties": {
      "containers": [
        {
          "command": [
            "/bin/sh"
          ],
          "args": [
            "-c",
            "echo hello world"
        }
      ]
    }
  }
}
E0F
```

```
$ aws batch submit-job --cli-input-json file://./submit-job-override.json
```

#### **Notes**

- For improved visibility into the details of the operations, enable Amazon EKS control plane logging. For more information, see <u>Amazon EKS control plane logging</u> in the <u>Amazon EKS User Guide</u>.
- Daemonsets and kubelets overhead affects available vCPU and memory resources, specifically scaling and job placement. For more information, see <u>Memory and vCPU considerations for</u> Amazon Batch on Amazon EKS.

To view the Job's output, do the following:

- Open the Amazon Batch console at <a href="https://console.amazonaws.cn/batch/">https://console.amazonaws.cn/batch/</a>.
- 2. In the navigation pane choose **Jobs**.
- 3. In the **Job queue** drop down choose the Job queue you created for the tutorial.
- 4. The **Jobs** table lists all of your Jobs and what their current status is. Once the Job's **Status** is **Succeeded** choose the **Name** of the Job to view the Job's details.
- 5. In the **Details** pane choose **Log stream name**. The CloudWatch console for the Job will open and there should be one event with the **Message** of hello world or your custom message.

# Step 9: Clean up your tutorial resources

You are charged for the Amazon EC2 instance while it is enabled. You can delete the instance to stop incurring charges.

To delete the resources you created, do the following:

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. In the navigation pane choose **Job queue**.
- 3. In the **Job queue** table choose the Job queue you created for the tutorial.
- 4. Choose **Disable**. Once the Job queue **State** is Disabled you can choose **Delete**.
- 5. Once the Job queue is deleted, in the navigation pane choose **Compute environments**.
- 6. Choose the compute environment you created for this tutorial and then choose **Disable**. It may take 1–2 minuets for the compute environment to complete being disabled.
- 7. Once the compute environment's **State** is Disabled, choose **Delete**. It may take 1–2 minuets for the compute environment to be deleted.

# **Additional resources**

After you complete the tutorial, you might want to explore the following topics::

- Learn more about the **Best practices**.
- Explore the Amazon Batch core components. For more information, see <u>Components of Amazon</u> Batch.

- Learn more about the different Compute Environments available in Amazon Batch.
- Learn more about Job queues and their different scheduling options.
- Learn more about Job definitions and the different configuration options.
- Learn more about the different types of <u>Jobs</u>.

# Getting started with Amazon Batch on Amazon EKS Private Clusters

Amazon Batch is a managed service that orchestrates batch workloads in your Amazon Elastic Kubernetes Service (Amazon EKS) clusters. This includes queuing, dependency tracking, managed job retries and priorities, pod management, and node scaling. This feature connects your existing private Amazon EKS cluster with Amazon Batch to run your jobs at scale. You can use <a href="mailto:eksctl">eksctl</a> (a command line interface for Amazon EKS), the Amazon console, or the <a href="mailto:Amazon Command Line">Amazon Command Line</a> Interface to create a private Amazon EKS cluster with all the other necessary resources.

Amazon EKS private only clusters by default have no inbound/outbound internet access, and you can only access the API server from within your VPC or a connected network. Amazon VPC endpoints are used to enable private access to other Amazon services. eksctl supports creating fully private clusters using a pre-existing Amazon VPC and subnets. eksctl also creates Amazon VPC endpoints in the supplied Amazon VPC and modifies route tables for the supplied subnets.

Each subnet should have an explicit route table associated with it because eksctl does not modify the main route table. Your <u>cluster</u> must pull images from a container registry that's in your Amazon VPC. As well, you can create an Amazon Elastic Container Registry in your Amazon VPC and copy container images to it for your nodes to pull from. For more information, see <u>Copy a container image from one repository to another repository</u>. To get started with Amazon ECR private repositories, see <u>Amazon ECR private repositories</u>.

You can optionally create a <u>pull through cache rule</u> with Amazon ECR. Once a pull through cache rule is created for an external public registry, you can pull an image from that external public registry using your Amazon ECR private registry uniform resource identifier (URI). Then Amazon ECR creates a repository and caches the image. When a cached image is pulled using the Amazon ECR private registry URI, Amazon ECR checks the remote registry to see if there is a new version of the image and updates your private registry up to once every 24 hours.

#### **Contents**

- Overview
- Prerequisites
- Step 1: Create your EKS cluster for Amazon Batch
- Step 2: Prepare your EKS cluster for Amazon Batch
- Step 3: Create an Amazon EKS compute environment
- Step 4: Create a job queue and attach the compute environment
- Step 5: Create an Amazon ECR with pull through cache
- Step 6: Register a job definition
- Step 7: Submit a job to run
- Step 8: View the Job's output
- Step 9: (Optional) Submit a job with overrides
- Step 10: Clean up your tutorial resources
- Additional resources
- Troubleshooting

#### **Overview**

This tutorial demonstrates how to setup Amazon Batch with a private Amazon EKS using the Amazon CloudShell, kubectl and eksctl.

#### **Intended Audience**

This tutorial is designed for system administrators and developers responsible for setting up, testing, and deploying Amazon Batch.

#### **Features Used**

This tutorial shows you how to use the Amazon CLI, to:

- Use Amazon Elastic Container Registry (Amazon ECR) to store container images
- Create and configure an Amazon EKS compute environment
- Create a job queue.
- Create a job definition
- Create and submit a job to run

Overview 49

Submit a job with overrides

#### **Time Required**

It should take about 40–50 minutes to complete this tutorial.

#### **Regional Restrictions**

There are no country or regional restrictions associated with using this solution.

#### **Resource Usage Costs**

There's no charge for creating an Amazon account. However, by implementing this solution, you might incur some or all of the costs that are listed in the following table.

Description	Cost (US dollars)
You are charged by the cluster hour	Varies depending on Instance, see <u>Amazon</u> <u>EKS pricing</u>
Amazon EC2 instance	You pay for each Amazon EC2 instance that is created. For more information about pricing, see <a href="Mazon EC2 Pricing"><u>Amazon EC2 Pricing</u></a> .

# **Prerequisites**

This tutorial uses Amazon CloudShell which is a browser-based, pre-authenticated shell that you launch directly from the Amazon Web Services Management Console. This allows for access to the cluster once it no longer has public internet access. The Amazon CLI, kubect1, and eksct1 may already be installed as part of Amazon CloudShell. For more information on Amazon CloudShell, see the <a href="Amazon CloudShellUser Guide">Amazon CloudShell User Guide</a>. An alternative to Amazon CloudShell is to connect to your cluster's VPC or a connected network.

To run kubectl commands, you will need private access to your Amazon EKS cluster. This means all traffic to your cluster API server must come from within your cluster's VPC or a connected network.

Amazon CLI – A command line tool for working with Amazon services, including Amazon EKS.
 This guide requires that you use version 2.8.6 or later or 1.26.0 or later. For more information, see <u>Installing, updating, and uninstalling the Amazon CLI</u> in the *Amazon Command Line Interface User Guide*. After installing the Amazon CLI, we recommend that you also configure it. For more

Prerequisites 50

information, see Quick configuration with aws configure in the Amazon Command Line Interface User Guide.

- kubect1 A command line tool for working with Kubernetes clusters. This guide requires that you use version 1.23 or later. For more information, see Installing or updating kubectl in the Amazon EKS User Guide.
- eksctl A command line tool for working with Amazon EKS clusters that automates many individual tasks. This guide requires that you use version 0.115.0 or later. For more information, see Installing or updating eksctl in the Amazon EKS User Guide.
- Permissions Users calling the CreateComputeEnvironment API operation to create a compute environment that uses Amazon EKS resources require permissions to the eks:DescribeCluster and eks:ListClusters API operation. You can attach the AWSBatchFullAccess managed policy to your user account by following the directions Adding and removing IAM identity permissions in the IAM User Guide.
- InstanceRole You need to create an InstanceRole for your Amazon EKS nodes that has the AmazonEKSWorkerNodePolicy and AmazonEC2ContainerRegistryPullOnly polices. For directions on how to create the InstanceRole see Creating the Amazon EKS node IAM role. You will need the ARN of the InstanceRole.
- Amazon Web Services account ID You need to know your Amazon Web Services account ID. Follow the directions in Viewing your Amazon Web Services account ID.
- (Optional) CloudWatch To examine the details of (Optional) Submit a job with overrides, logging must be configured. For more information, see Use CloudWatch Logs to monitor Amazon Batch on Amazon EKS jobs.

# Step 1: Create your EKS cluster for Amazon Batch



#### Important

To get started as simply and quickly as possible, this tutorial includes steps with default settings. Before creating for production use, we recommend that you familiarize yourself with all settings and deploy with the settings that meet your requirements.

We recommend you use eksctl and the following config file to create your cluster. To manually setup your cluster follow the directions in Deploy private clusters with limited internet access in the Amazon EKS User Guide.

1. Open the <u>Amazon CloudShell console</u> and set the region to us-east-1. For the rest of the tutorial make sure you are using us-east-1.

Create a private EKS cluster in the us-east-1 region using the sample eksctl
config file. Save the yaml file to your Amazon CloudShell environment and name it
clusterConfig.yaml. You can change my-test-cluster with the name you want to use
for your cluster.

```
kind: ClusterConfig
apiVersion: eksctl.io/v1alpha5
metadata:
    name: my-test-cluster
    region: us-east-1
availabilityZones:
    - us-east-1a
    - us-east-1b
    - us-east-1c
managedNodeGroups:
    - name: ng-1
    privateNetworking: true
privateCluster:
    enabled: true
    skipEndpointCreation: false
```

- Create your resources using the command: eksctl create cluster -f clusterConfig.yaml. Cluster creation can take between 10–15 minutes.
- 4. Once the cluster has finished being created you have to add your Amazon CloudShell IP address to the allow list. To find your Amazon CloudShell IP address run the following command:

```
curl http://checkip.amazonaws.com
```

Once you have the public IP address you have to create an allow list rule:

```
aws eks update-cluster-config \
    --name my-test-cluster \
    --region us-east-1 \
    --resources-vpc-config
endpointPublicAccess=true,endpointPrivateAccess=true,publicAccessCidrs=["<Public IP>/32"]
```

Then apply the update to the kubectl config file:

```
aws eks update-kubeconfig --name my-test-cluster --region us-east-1
```

5. To test that you have access to the nodes run the following command:

```
kubectl get nodes
```

The output of command is:

```
NAME
                                   STATUS
                                            ROLES
                                                     AGE
                                                             VERSION
ip-192-168-107-235.ec2.internal
                                  Ready
                                                     1h
                                                          v1.32.3-eks-473151a
                                            none
ip-192-168-165-40.ec2.internal
                                  Ready
                                                     1h
                                                          v1.32.3-eks-473151a
                                            none
ip-192-168-98-54.ec2.internal
                                                          v1.32.1-eks-5d632ec
                                  Ready
                                                     1h
                                            none
```

# Step 2: Prepare your EKS cluster for Amazon Batch

All steps are required and must be done in Amazon CloudShell.

1. Create a dedicated namespace for Amazon Batch jobs

Use kubectl to create a new namespace.

```
$ namespace=my-aws-batch-namespace
```

```
$ cat - <<EOF | kubectl create -f -
{
  "apiVersion": "v1",
  "kind": "Namespace",
  "metadata": {
     "name": "${namespace}",
     "labels": {
        "name": "${namespace}"
     }
  }
}
EOF</pre>
```

Output:

namespace/my-aws-batch-namespace created

#### 2. Enable access via role-based access control (RBAC)

Use kubectl to create a Kubernetes role for the cluster to allow Amazon Batch to watch nodes and pods, and to bind the role. You must do this once for each Amazon EKS cluster.

```
$ cat - <<EOF | kubectl apply -f -</pre>
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: aws-batch-cluster-role
rules:
  - apiGroups: [""]
    resources: ["namespaces"]
    verbs: ["get"]
  - apiGroups: [""]
    resources: ["nodes"]
    verbs: ["get", "list", "watch"]
  - apiGroups: [""]
    resources: ["pods"]
    verbs: ["get", "list", "watch"]
  - apiGroups: [""]
    resources: ["events"]
    verbs: ["list"]
  - apiGroups: [""]
    resources: ["configmaps"]
    verbs: ["get", "list", "watch"]
  - apiGroups: ["apps"]
    resources: ["daemonsets", "deployments", "statefulsets", "replicasets"]
    verbs: ["get", "list", "watch"]
  - apiGroups: ["rbac.authorization.k8s.io"]
    resources: ["clusterroles", "clusterrolebindings"]
    verbs: ["get", "list"]
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: aws-batch-cluster-role-binding
subjects:
- kind: User
  name: aws-batch
```

```
apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: ClusterRole
  name: aws-batch-cluster-role
  apiGroup: rbac.authorization.k8s.io
EOF
```

#### Output:

```
clusterrole.rbac.authorization.k8s.io/aws-batch-cluster-role created clusterrolebinding.rbac.authorization.k8s.io/aws-batch-cluster-role-binding created
```

Create namespace-scoped Kubernetes role for Amazon Batch to manage and lifecycle pods and bind it. You must do this once for each unique namespace.

```
$ namespace=my-aws-batch-namespace
```

```
$ cat - <<EOF | kubectl apply -f - --namespace "${namespace}"</pre>
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: aws-batch-compute-environment-role
  namespace: ${namespace}
rules:
  - apiGroups: [""]
    resources: ["pods"]
    verbs: ["create", "get", "list", "watch", "delete", "patch"]
  - apiGroups: [""]
    resources: ["serviceaccounts"]
    verbs: ["get", "list"]
  - apiGroups: ["rbac.authorization.k8s.io"]
    resources: ["roles", "rolebindings"]
    verbs: ["get", "list"]
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: aws-batch-compute-environment-role-binding
  namespace: ${namespace}
subjects:
- kind: User
```

```
name: aws-batch
apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: Role
  name: aws-batch-compute-environment-role
  apiGroup: rbac.authorization.k8s.io
EOF
```

#### Output:

```
role.rbac.authorization.k8s.io/aws-batch-compute-environment-role created
rolebinding.rbac.authorization.k8s.io/aws-batch-compute-environment-role-binding
  created
```

Update Kubernetes aws-auth configuration map to map the preceding RBAC permissions to the Amazon Batch service-linked role.

```
$ eksctl create iamidentitymapping \
    --cluster my-test-cluster \
    --arn "arn:aws-cn:iam::<your-account-ID>:role/AWSServiceRoleForBatch" \
    --username aws-batch
```

#### Output:

```
2022-10-25 20:19:57 [#] adding identity "arn:aws-cn:iam::
vour-account-ID>: role/AWSServiceRoleForBatch" to auth ConfigMap
```

# Note

The path aws-service-role/batch.amazonaws.com/ has been removed from the ARN of the service-linked role. This is because of an issue with the aws-auth configuration map. For more information, see Roles with paths don't work when the path is included in their ARN in the aws-authconfigmap.

# Step 3: Create an Amazon EKS compute environment

Amazon Batch compute environments define compute resource parameters to meet your batch workload needs. In a managed compute environment, Amazon Batch helps you to manage the

capacity and instance types of the compute resources (Kubernetes nodes) within your Amazon EKS cluster. This is based on the compute resource specification that you define when you create the compute environment. You can use EC2 On-Demand Instances or EC2 Spot Instances.

Now that the **AWSServiceRoleForBatch** service-linked role has access to your Amazon EKS cluster, you can create Amazon Batch resources. First, create a compute environment that points to your Amazon EKS cluster.

- For subnets run eksctl get cluster *my-test-cluster* to get the subnets used by the cluster.
- For securityGroupIds parameter you can use the same security group as the Amazon EKS cluster. This command retrieves the security group ID for the cluster.

```
$ aws eks describe-cluster \
    --name my-test-cluster \
    --query cluster.resourcesVpcConfig.clusterSecurityGroupId
```

• Use the ARN of the instanceRole you created in the Prerequisites.

```
$ cat <<EOF > ./batch-eks-compute-environment.json
{
  "computeEnvironmentName": "My-Eks-CE1",
  "type": "MANAGED",
  "state": "ENABLED",
  "eksConfiguration": {
    "eksClusterArn": "arn:aws-cn:eks:us-east-1:<your-account-ID>:cluster/my-test-
cluster",
    "kubernetesNamespace": "my-aws-batch-namespace"
  },
  "computeResources": {
    "type": "EC2",
    "allocationStrategy": "BEST_FIT_PROGRESSIVE",
    "minvCpus": 0,
    "maxvCpus": 128,
    "instanceTypes": [
        "m5"
    ],
    "subnets": [
        "<eks-cluster-subnets-with-access-to-the-image-for-image-pull>"
    "securityGroupIds": [
```

```
"<eks-cluster-sg>"
],
"instanceRole": "<eks-instance-profile>"
}
EOF
```

```
\$ aws batch create-compute-environment --cli-input-json file://./batch-eks-compute-environment.json
```

#### **Notes**

• Maintenance of an Amazon EKS compute environment is a shared responsibility. For more information, see Security in Amazon EKS.

# Step 4: Create a job queue and attach the compute environment

# ∧ Important

It's important to confirm that the compute environment is healthy before proceeding. The DescribeComputeEnvironments API operation can be used to do this.

```
$ aws batch describe-compute-environments --compute-environments My-Eks-CE1
```

Confirm that the status parameter is not INVALID. If it is, look at the statusReason parameter for the cause. For more information, see <u>Troubleshooting Amazon Batch</u>.

Jobs submitted to this new job queue are run as pods on Amazon Batch managed nodes that joined the Amazon EKS cluster that's associated with your compute environment.

```
$ cat <<EOF > ./batch-eks-job-queue.json
{
    "jobQueueName": "My-Eks-JQ1",
    "priority": 10,
    "computeEnvironmentOrder": [
        {
            "order": 1,
            "order": 1,
```

```
$ aws batch create-job-queue --cli-input-json file://./batch-eks-job-queue.json
```

# Step 5: Create an Amazon ECR with pull through cache

Because the cluster doesn't have public internet access you have to create an Amazon ECR for container images. The following directions create an Amazon ECR with a pull-through cache rule to store the image.

 The following command create the pull-through cache rule. You can replace tutorialprefix with a different prefix.

```
aws ecr create-pull-through-cache-rule \
    --ecr-repository-prefix "my-prefix" \
    --upstream-registry-url "public.ecr.aws" \
    --region us-east-1
```

2. Authenticate with the public ECR.

```
aws ecr get-login-password --region us-east-1 | docker login --username AWS --password-stdin <your-account-ID>.dkr.ecr.us-east-1.amazonaws.com
```

Now you can pull an image.

```
docker pull <your-account-ID>.dkr.ecr.us-east-1.amazonaws.com/my-prefix/
amazonlinux/amazonlinux:2
```

3. You can verify the repositry and image by running the following commands:

```
aws ecr describe-repositories

aws ecr describe-images --repository-name my-prefix/amazonlinux/amazonlinux
```

4. The image string to use for pulling the container is in the following format:

```
<your-account-ID>.dkr.ecr.us-east-1.amazonaws.com/my-prefix/amazonlinux/
amazonlinux:2
```

# Step 6: Register a job definition

The following Job definition instructs the pod to sleep for 60 seconds.

In the image field of the job definition, instead of providing a link to image in a public ECR repository, provide the link to the image stored in our private ECR repository. See the following sample job definition:

```
$ cat <<EOF > ./batch-eks-job-definition.json
  "jobDefinitionName": "MyJobOnEks_Sleep",
  "type": "container",
  "eksProperties": {
    "podProperties": {
      "hostNetwork": true,
      "containers": [
          "image": "<your-account-ID>.dkr.ecr.us-east-1.amazonaws.com/my-prefix/
amazonlinux/amazonlinux:2",
          "command": [
            "sleep",
            "60"
          ],
          "resources": {
            "limits": {
              "cpu": "1",
              "memory": "1024Mi"
            }
          }
        }
      ],
      "metadata": {
        "labels": {
          "environment": "test"
        }
      }
    }
  }
```

```
}
EOF
```

```
$ aws batch register-job-definition --cli-input-json file://./batch-eks-job-
definition.json
```

### **Notes**

There are considerations for the cpu and memory parameters. For more information, see
 Memory and vCPU considerations for Amazon Batch on Amazon EKS.

## Step 7: Submit a job to run

Run the following Amazon CLI command in Amazon CloudShell to submit a new Job and returns the unique JobID.

```
$ aws batch submit-job --job-queue My-Eks-JQ1 \
    --job-definition MyJobOnEks_Sleep - -job-name My-Eks-Job1
```

### **Notes**

• For more information about running jobs on Amazon EKS resources, see Amazon EKS jobs.

## Step 8: View the Job's output

To check the status of a Job:

```
$ aws batch describe-jobs --job <JobID-from-submit-response>
```

The startedAt and stoppedAt should be one minute apart.

# Step 9: (Optional) Submit a job with overrides

This job overrides the command passed to the container.

```
$ cat <<EOF > ./submit-job-override.json
{
   "jobName": "EksWithOverrides",
```

Step 7: Submit a job to run 61

```
"jobQueue": "My-Eks-JQ1",
  "jobDefinition": "MyJobOnEks_Sleep",
  "eksPropertiesOverride": {
    "podProperties": {
      "containers": [
        {
          "command": [
             "/bin/sh"
          ],
          "args": [
            "-c",
            "echo hello world"
          ]
      ]
    }
  }
}
EOF
```

```
$ aws batch submit-job - -cli-input-json file://./submit-job-override.json
```

### **Notes**

- For improved visibility into the details of the operations, enable Amazon EKS control plane logging. For more information, see <u>Amazon EKS control plane logging</u> in the Amazon EKS User Guide.
- Daemonsets and kubelets overhead affects available vCPU and memory resources, specifically scaling and job placement. For more information, see <u>Memory and vCPU considerations for</u> Amazon Batch on Amazon EKS.

# Step 10: Clean up your tutorial resources

You are charged for the Amazon EC2 instance while it is enabled. You can delete the instance to stop incurring charges.

To delete the resources you created, do the following:

1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.

- 2. In the navigation pane choose **Job queue**.
- 3. In the **Job queue** table choose the Job queue you created for the tutorial.
- 4. From **Actions** choose **Disable**. Once the Job queue **State** is Disabled you can choose **Delete**.
- 5. Once the Job queue is deleted, in the navigation pane choose **Compute environments**.
- 6. Choose the compute environment you created for this tutorial and then choose **Disable** from **Actions**. It may take 1–2 minuets for the compute environment to complete being disabled.
- Once the compute environment's State is Disabled, choose Delete. It may take 1–2 minuets for the compute environment to be deleted.

### **Additional resources**

After you complete the tutorial, you might want to explore the following topics::

- · Learn more about the Best practices.
- Explore the Amazon Batch core components. For more information, see <a href="Components of Amazon">Components of Amazon</a>
  Batch.
- Learn more about the different Compute Environments available in Amazon Batch.
- Learn more about Job queues and their different scheduling options.
- Learn more about <u>Job definitions</u> and the different configuration options.
- Learn more about the different types of Jobs.

# **Troubleshooting**

If nodes launched by Amazon Batch don't have access to the Amazon ECR repository (or any other repository) that stores your image, then your jobs could remain in the STARTING state. This is because the pod will not be able to download the image and run your Amazon Batch job. If you click on the pod name launched by Amazon Batch you should be able to see the error message and confirm the issue. The error message should look similar to the following:

```
Failed to pull image "public.ecr.aws/amazonlinux/amazonlinux:2": rpc error: code = Unknown desc = failed to pull and unpack image "public.ecr.aws/amazonlinux/amazonlinux:2": failed to resolve reference "public.ecr.aws/amazonlinux/amazonlinux:2": failed to do request: Head "https://public.ecr.aws/v2/amazonlinux/amazonlinux/manifests/2": dial tcp: i/o timeout
```

Additional resources 63

For other common troubleshooting scenarios, see <u>Troubleshooting Amazon Batch</u>. For troubleshooting based on pod status, see <u>How do I troubleshoot the pod status in Amazon EKS?</u>.

Troubleshooting 64

# The Amazon Batch widget dashboard

Through the Amazon Batch dashboard, you can monitor recent jobs, job queues, and compute environments. By default, the following dashboard widgets are displayed:

- Job overview For more information about Amazon Batch jobs, see Jobs.
- Job queue overview For more information about Amazon Batch job queues, see Job queues.
- **Compute environment overview** For more information about Amazon Batch compute environments, see *Compute environments for Amazon Batch*.

You can customize the widgets that are displayed on the Dashboard page. The following sections describe additional widgets that you can add.

### **Topics**

- · How to add the Single job queue widget to the Amazon Batch dashboard
- How to add the CloudWatch Container Insights widget to the Amazon Batch dashboard
- How to add the Job logs widget to the Amazon Batch dashboard

# How to add the Single job queue widget to the Amazon Batch dashboard

The Single job queue widget displays detailed information about a single job queue.

To add this widget, follow these steps.

- 1. Open the Amazon Batch console.
- 2. From the navigation bar, choose the Amazon Web Services Region that you want.
- 3. In the navigation pane, choose **Dashboard**.
- 4. Choose **Add widgets**.
- 5. For Single job queue, choose Add widget.
- 6. For **Job queue**, select the job queue that you want.
- 7. For **Job status**, choose the job statuses that you want to display.
- 8. (Optional) Turn off **Show connected compute environments** if you don't want to display the properties for compute environments.

9. For **Compute environment properties**, select the properties that you want.

10Choose Add.

# How to add the CloudWatch Container Insights widget to the Amazon Batch dashboard

This widget displays aggregated metrics for Amazon Batch compute environments and jobs. For more information about Container Insights, see the section called "CloudWatch Container Insights".

To add this widget, follow these steps.

- 1. Open the Amazon Batch console.
- 2. From the navigation bar, choose the Amazon Web Services Region that you want.
- 3. In the navigation pane, choose **Dashboard**.
- 4. Choose Add widgets.
- 5. For Container insights, choose Add widget.
- 6. For **Compute environment**, choose the compute environment that you want.
- 7. Choose Add.

# How to add the Job logs widget to the Amazon Batch dashboard

This widget displays different logs from your jobs in one convenient location. For more information about job logs, see the section called "Tutorial: Send Amazon Batch job logs to CloudWatch Logs".

To add this widget, follow these steps.

- 1. Open the Amazon Batch console.
- 2. From the navigation bar, choose the Amazon Web Services Region that you want.
- 3. In the navigation pane, choose **Dashboard**.
- 4. Choose Add widgets.
- 5. For **Job logs**, choose **Add widget**.
- 6. For **Job id**, enter the job ID for the job that you want.

### 7. Choose **Add**.

Add the Job logs widget 67

# **Compute environments for Amazon Batch**

Job queues are mapped to one or more compute environments. Compute environments contain the Amazon ECS container instances that are used to run containerized batch jobs. A specific compute environment can also be mapped to one or more than one job queue. Within a job queue, the associated compute environments each have an order that's used by the scheduler to determine where jobs that are ready to be run will run. If the first compute environment has a status of VALID and has available resources, the job is scheduled to a container instance within that compute environment. If the first compute environment has a status of INVALID or can't provide a suitable compute resource, the scheduler attempts to run the job on the next compute environment.

### **Topics**

- Managed compute environments
- Unmanaged compute environments
- Create a compute environment
- Compute resource AMIs
- Use Amazon EC2 launch templates with Amazon Batch
- EC2 configurations
- Instance type allocation strategies for Amazon Batch
- Compute resource memory management
- Updating compute environments
- Fargate compute environments
- Amazon EKS compute environments

# Managed compute environments

You can use a managed compute environment to have Amazon Batch manage the capacity and instance types of the compute resources within the environment. This is based on the compute resource specifications that you define when you create the compute environment. You can choose either to use Amazon EC2 On-Demand Instances and Amazon EC2 Spot Instances. Or, you can alternatively use Fargate and Fargate Spot capacity in your managed compute environment. When using Spot Instances, you can optionally set a maximum price. This way, Spot Instances only launch when the Spot Instance price is under a specified percentage of the On-Demand price.

### Important

Fargate Spot instances are not supported on Windows containers on Amazon Fargate. A job queue will be blocked if a FargateWindows job is submitted to a job queue that only uses Fargate Spot compute environments.

### Important

Amazon Batch creates and manages multiple Amazon resources on your behalf and within your account, including Amazon EC2 Launch Templates, Amazon EC2 Auto Scaling Groups, Amazon EC2 Spot Fleets, and Amazon ECS Clusters. These managed resources are configured specifically to ensure optimal Amazon Batch operation. Manually modifying these Batch-managed resources, unless explicitly stated in Amazon Batch documentation, may result in unexpected behavior resulting in INVALID Compute Environment, suboptimal instance scaling behavior, delayed workload processing, or unexpected costs. These manual modifications can not be deterministically supported by the Amazon Batch service. Always use the supported Batch APIs or the Batch console to manage your Compute Environments.

Managed compute environments launch Amazon EC2 instances into the VPC and subnets that you specify and then registers them with an Amazon ECS cluster. The Amazon EC2 instances need external network access to communicate with the Amazon ECS service endpoint. Some subnets don't provide Amazon EC2 instances with public IP addresses. If your Amazon EC2 instances don't have a public IP address, they must use network address translation (NAT) to gain this access. For more information, see NAT gateways in the Amazon VPC User Guide. For more information about how to create a VPC, see Create a virtual private cloud.

By default, Amazon Batch managed compute environments use a recent, approved version of the Amazon ECS optimized AMI for compute resources. However, you might want to create your own AMI to use for your managed compute environments for various reasons. For more information, see Compute resource AMIs.



### Note

Amazon Batch doesn't automatically upgrade the AMIs in a compute environment after it's created. For example, it doesn't update the AMIs in your compute environment when

a newer version of the Amazon ECS optimized AMI is released. You're responsible for the management of the guest operating system. This includes any updates and security patches. You're also responsible for any additional application software or utilities that you install on the compute resources. There are two ways to use a new AMI for your Amazon Batch jobs. The original method is to complete these steps:

- 1. Create a new compute environment with the new AMI.
- 2. Add the compute environment to an existing job queue.
- 3. Remove the earlier compute environment from your job queue.
- 4. Delete the earlier compute environment.

In April 2022, Amazon Batch added enhanced support for updating compute environments. For more information, see <u>Updating compute environments</u>. To use the enhanced updating of compute environments to update AMIs, follow these rules:

- Either don't set the service role (<u>serviceRole</u>) parameter or set it to the **AWSServiceRoleForBatch** service-linked role.
- Set the allocation strategy (<u>allocationStrategy</u>) parameter to BEST\_FIT\_PROGRESSIVE, SPOT\_CAPACITY\_OPTIMIZED or SPOT\_PRICE\_CAPACITY\_OPTIMIZED.
- Set the update to latest image version (<u>updateToLatestImageVersion</u>) parameter to true.
- Don't specify an AMI ID in <a href="imageId">imageId</a>, <a href="imageId">imageId</a> (in <a href="ec2Configuration">ec2Configuration</a>), or in the launch template (<a href="launchTemplate">launchTemplate</a>). In that case, Amazon Batch selects the latest Amazon ECS optimized AMI that's supported by Amazon Batch at the time the infrastructure update is initiated. Alternatively, you can specify the AMI ID in the <a href="imageId">imageId</a> or <a href

If these rules are followed, any update that starts an infrastructure update will cause the AMI ID to be re-selected. If the version setting in the launch template (launchTemplate) is set to \$Latest or \$Default, the latest or default version of the launch template are evaluated up at the time of the infrastructure update, even if the launchTemplate was not updated.

# Consideration when creating multi-node parallel jobs

Amazon Batch recommends creating dedicated compute environments for running multi-node parallel (MNP) jobs and non-MNP jobs. This is due to the way compute capacity is created in your managed compute environment. When creating a new managed compute environment, if you specify a minvCpu value greater than zero then Amazon Batch creates an instance pool for use with non-MNP jobs only. If a multi-node parallel job is submitted, Amazon Batch creates new instance capacity to run the multi-node parallel jobs. In cases where there are both single-node and multi-node parallel jobs running in the same compute environment where either a minvCpus or maxvCpus value is set, if the required compute resources are unavailable Amazon Batch will wait for the current jobs to finish before creating the compute resources necessary to run the new jobs.

# **Unmanaged compute environments**

In an unmanaged compute environment, you manage your own compute resources. You must verify that the AMI you use for your compute resources meets the Amazon ECS container instance AMI specification. For more information, see Compute resource AMI specification and Tutorial: Create a compute resource AMI.



### Note

Amazon Fargate resources aren't supported in unmanaged compute environments.

After you created your unmanaged compute environment, use the DescribeComputeEnvironments API operation to view the compute environment details. Find the Amazon ECS cluster that's associated with the environment and then manually launch your container instances into that Amazon ECS cluster.

The following Amazon CLI command also provides the Amazon ECS cluster ARN.

```
$ aws batch describe-compute-environments \
    --compute-environments unmanagedCE \
    --query "computeEnvironments[].ecsClusterArn"
```

For more information, see <u>Launching an Amazon ECS container instance</u> in the *Amazon Elastic Container Service Developer Guide*. When you launch your compute resources, specify the Amazon ECS cluster ARN that the resources register with the following Amazon EC2 user data. Replace <u>ecsClusterArn</u> with the cluster ARN that you obtained with the previous command.

```
#!/bin/bash
echo "ECS_CLUSTER=ecsClusterArn" >> /etc/ecs/ecs.config
```

# Create a compute environment

Before you can run jobs in Amazon Batch, you need to create a compute environment. You can create a managed compute environment where Amazon Batch manages the Amazon EC2 instances or Amazon Fargate resources within the environment based on your specifications. Or, alternatively, you can create an unmanaged compute environment where you handle the Amazon EC2 instance configuration within the environment.

### 

Fargate Spot instances are not supported in the following scenarios:

- On Amazon Linux containers with ARM64 architecture.
- Windows containers on Amazon Fargate

A job queue will be blocked in these scenarios if a job is submitted to a job queue that only uses Fargate Spot compute environments.

### **Topics**

- Tutorial: Create a managed compute environment using Fargate resources
- Tutorial: Create a managed compute environment using Amazon EC2 resources
- Tutorial: Create an unmanaged compute environment using Amazon EC2 resources
- Tutorial: Create a managed compute environment using Amazon EKS resources

Resource: Compute environment template

# Tutorial: Create a managed compute environment using Fargate resources

Complete the following steps to create a managed compute environment using Amazon Fargate resources.

- Open the Amazon Batch console at https://console.amazonaws.cn/batch/. 1.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Compute environments**.
- 4. Choose Create.
- 5. Configure the compute environment.



### Note

Compute environments for Windows containers on Amazon Fargate jobs must at least one vCPU.

- For Compute environment configuration, choose Fargate.
- For Name, specify a unique name for your compute environment. The name can contain b. up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- For **Service role**, choose service-linked role that lets the Amazon Batch service to make calls to the required Amazon API operations on your behalf. For example, choose AWSServiceRoleForBatch. For more information, see Service-linked role permissions for Amazon Batch.
- (Optional) Expand Tags. To add a tag, choose Add tag. Then, enter a Key name and optional Value. Choose Add tag.
- Choose **Next page**.
- 6. In the **Instance configuration** section:
  - (Optional) For Use Fargate Spot capacity, turn on Fargate Spot. For information about Fargate Spot, see Using Amazon EC2 Spot and Fargate\_SPOT.

For Maximum vCPUs, choose the maximum number of vCPUs that your compute environment can scale out to, regardless of job gueue demand.

- Choose **Next page**. c.
- Configure networking. 7.

### Important

Compute resources need access to communicate with the Amazon ECS service endpoint. This can be through an interface VPC endpoint or through your compute resources having public IP addresses.

For more information about interface VPC endpoints, see Amazon ECS Interface VPC Endpoints (Amazon PrivateLink) in the Amazon Elastic Container Service Developer Guide.

If you do not have an interface VPC endpoint configured and your compute resources do not have public IP addresses, then they must use network address translation (NAT) to provide this access. For more information, see NAT gateways in the Amazon VPC User Guide. For more information, see the section called "Create a VPC".

- For Virtual Private Cloud (VPC) ID, choose a VPC where you want to launch your a. instances.
- For **Subnets**, choose the subnets to use. By default, all subnets within the selected VPC are available.



### Note

Amazon Batch on Fargate doesn't currently support Local Zones. For more information, see Amazon ECS clusters in Local Zones, Wavelength Zones, and Amazon Outposts in the Amazon Elastic Container Service Developer Guide.

- For **Security groups**, choose a security group to attach to your instances. By default, the C. default security group for your VPC is chosen.
- Choose **Next page**. d.
- For **Review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create compute environment**.

# Tutorial: Create a managed compute environment using Amazon EC2 resources

Complete the following steps to create a managed compute environment using Amazon Elastic Compute Cloud (Amazon EC2) resources.

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Compute environments**.
- 4. Choose Create.
- 5. Configure the environment.
  - a. For Compute environment configuration, choose Amazon Elastic Compute Cloud (Amazon EC2).
  - b. For Orchestration type, choose Managed.
  - c. For **Name**, specify a unique name for your compute environment. The name can contain up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
  - d. For **Service role**, choose service-linked role that lets the Amazon Batch service make calls to the required Amazon API operations on your behalf. For example, choose **AWSServiceRoleForBatch**. For more information, see <u>Service-linked role permissions for Amazon Batch</u>.
  - e. For **Instance role**, choose to create a new instance profile or use an existing instance profile that has the required IAM permissions attached. This instance profile allows the Amazon ECS container instances that are created for your compute environment to make calls to the required Amazon API operations on your behalf. For more information, see <a href="Amazon ECS instance role">Amazon ECS instance role</a>. If you choose to create a new instance profile, the required role (ecsInstanceRole) is created for you.
  - f. (Optional) Expand Tags.
  - g. (Optional) For **EC2 tags**, choose **Add tag** to add a tag to resources that are launched in the compute environment. Then, enter a **Key** name and optional **Value**. Choose **Add tag**.
  - h. (Optional) For **Tags**, choose **Add tag**. Then, enter a **Key** name and optional **Value**. Choose **Add tag**.

For more information, see Tag your Amazon Batch resources.

- i. Choose **Next page**.
- 6. In the **Instance configuration** section:
  - (Optional) For **Enable using Spot instances**, turn on Spot. For more information, see Spot a. Instances.
  - (Spot only) For Maximum % on-demand price, choose the maximum percentage that a Spot Instance price can be when compared with the On-Demand price for that instance type before instances are launched. For example, if your maximum price is 20%, then the Spot price must be less than 20% of the current On-Demand price for that EC2 instance. You always pay the lowest (market) price and never more than your maximum percentage. If you leave this field empty, the default value is 100% of the On-Demand price.
  - (Spot only) For **Spot fleet role**, choose an existing Amazon EC2 Spot Fleet IAM role to c. apply to your Spot compute environment. If you don't already have an existing Amazon EC2 Spot Fleet IAM role, you must create one first. For more information, see Amazon EC2 spot fleet role.

### Important

To tag your Spot Instances on creation, your Amazon EC2 Spot Fleet IAM role must use the newer **AmazonEC2SpotFleetTaggingRole** managed policy. The **AmazonEC2SpotFleetRole** managed policy doesn't have the required permissions to tag Spot Instances. For more information, see Spot Instances not tagged on creation and the section called "Tag your resources".

- d. For **Minimum vCPUs**, choose the minimum number of vCPUs that your compute environment maintains, regardless of job queue demand.
- For **Desired vCPUs**, choose the number of vCPUs that your compute environment launches with. As your job queue demand increases, Amazon Batch can increase the desired number of vCPUs in your compute environment and add EC2 instances, up to the maximum vCPUs. As demand decreases, Amazon Batch can decrease the desired number of vCPUs in your compute environment and remove instances, down to the minimum vCPUs.
- f. For **Maximum vCPUs**, choose the maximum number of vCPUs that your compute environment can scale out to, regardless of job queue demand.
- For **Allowed instance types**, choose the Amazon EC2 instance types that can be launched. You can specify instance families to launch any instance type within those families

(for example, c5, c5n, or p3). Or, you can specify specific sizes within a family (such as c5.8xlarge). Metal instance types aren't in the instance families. For example, c5 doesn't include c5.metal. You can also choose optimal to select instance types (from the C4, M4, and R4 instance families) that match the demand of your job gueues.

### (i) Note

When you create a compute environment, the instance types that you select for the compute environment must share the same architecture. For example, you can't mix x86 and ARM instances in the same compute environment.

### (i) Note

Amazon Batch will scale GPUs based on the required amount in your job queues. To use GPU scheduling, the compute environment must include instance types from the p3, p4, p5, p6, g3, g3s, g4, g5, or g6 families.

### Note

Currently, optimal uses instance types from the C4, M4, and R4 instance families. In Amazon Web Services Regions that don't have instance types from those instance families, instance types from the C5, M5, and R5 instance families are used.

- Expand **Additional configuration**. h.
- i. (Optional) For **Placement group**, enter a placement group name to group resources in the compute environment.
- (Optional) For **EC2 key pair**, choose a public and private key pair as security credentials j. when you connect to the instance. For more information about Amazon EC2 key pairs, see Amazon EC2 key pairs and Linux instances.
- For Allocation strategy, choose the allocation strategy to use when selecting instance types from the list of allowed instance types. **BEST\_FIT\_PROGRESSIVE** is usually the better choice for EC2 On-Demand compute environments, SPOT\_CAPACITY\_OPTIMIZED,

and SPOT\_PRICE\_CAPACITY\_OPTIMIZED for EC2 Spot compute environments. For more information, see the section called "Instance type allocation strategies".

l. (Optional) For **EC2 configuration** choose **Image type** and **Image ID override** values to provide information for Amazon Batch to select Amazon Machine Images (AMIs) for instances in the compute environment. If the **Image ID override** isn't specified for each Image type, Amazon Batch selects a recent Amazon ECS optimized AMI. If no Image type is specified, the default is a **Amazon Linux 2** for non-GPU, non Amazon Graviton instance.



### Important

To use a custom AMI, choose the image type and then enter the custom AMI ID in the **Image ID override** box.

### Amazon Linux 2

Default for all Amazon Graviton-based instance families (for example, C6g, M6g, R6g, and T4g) and can be used for all non-GPU instance types.

### Amazon Linux 2 (GPU)

Default for all GPU instance families (for example P4 and G4) and can be used for all non Amazon Graviton-based instance types.

### **Amazon Linux**

Can be used for non-GPU, non Amazon Graviton instance families. The standard support for Amazon Linux AMI has ended. For more information, see Amazon Linux AMI.



### Note

The AMI that you choose for a compute environment must match the architecture of the instance types that you intend to use for that compute environment. For example, if your compute environment uses A1 instance types, the compute resource AMI that you choose must support ARM instances. Amazon ECS vends both x86 and ARM versions of the Amazon ECS optimized Amazon Linux 2 AMI.

> For more information, see Amazon ECS optimized Amazon Linux 2 AMI in the Amazon Elastic Container Service Developer Guide.

m. (Optional) For Launch template, select an existing Amazon EC2 launch template to configure your compute resources. The default version of the template is automatically populated. For more information, see Use Amazon EC2 launch templates with Amazon Batch.



### Note

In a launch template, you can specify a custom AMI that you created.

(Optional) For Launch template version, enter \$Default, \$Latest, or a specific version n. number to use.



### Note

Note: If you use either substitution variable (\$Default or \$Latest), they will apply the current default or latest version number at the time that this configuration is saved. If the default or latest version changes in the future, you must update the information - it won't automatically update.

### 

If the version parameter of the launch template is \$Default or \$Latest, the default or latest version of the specified launch template is evaluated during an infrastructure update. If a different AMI ID is selected by the default or the latest version of the launch template is selected, that AMI ID is used in the update. For more information, see the section called "Updating the AMI ID".

- (Optional) For Override launch templates, select an existing Amazon EC2 launch template to use for specific instance types and families.
- (Optional) For Override launch template version, enter a specific version number to use, \$Default, or \$Latest.



### Note

If you use either the \$Default or \$Latest variable, Amazon Batch will apply the current information at the time that the compute environment is created. If the default or latest version changes in the future, you must update the information through UpdateComputeEnvironment or through the Amazon Web Services Management Console - Amazon Batch.

(Optional) For **Target instance types**, select the instance type or family that you want to apply the override launch template.



### Note

If you specify an override launch template, **Target instance types** is required. For more information, see LaunchTemplateSpecificationOverride.targetInstanceTypes.

### Note

If the instance type or family that you want to select doesn't appear in this list, review the selections you made in Allowed instance types.

- Choose **Next page**.
- In the **Network configuration** section: 7.

### Important

Compute resources need access to communicate with the Amazon ECS service endpoint. This can be through an interface VPC endpoint or through your compute resources having public IP addresses.

For more information about interface VPC endpoints, see Amazon ECS Interface VPC Endpoints (Amazon PrivateLink) in the Amazon Elastic Container Service Developer Guide.

If you do not have an interface VPC endpoint configured and your compute resources do not have public IP addresses, then they must use network address translation (NAT)

to provide this access. For more information, see NAT gateways in the Amazon VPC User Guide. For more information, see the section called "Create a VPC".

- For Virtual Private Cloud (VPC) ID, choose a VPC where to launch your instances. a.
- For **Subnets**, choose the subnets to use. By default, all subnets within the selected VPC are b. available.



### Note

Amazon Batch on Amazon EC2 supports Local Zones. For more information, see Local Zones in the Amazon EC2 User Guide and Amazon ECS clusters in Local Zones, Wavelength Zones, and Amazon Outposts in the Amazon Elastic Container Service Developer Guide.

(Optional) For **Security groups**, choose a security group to attach to your instances. By C. default, the default security group for your VPC is chosen.



### Note

Note: If you use either substitution variable (\$Default or \$Latest), they will apply the current default or latest version number at the time that this configuration is saved. If the default or latest version changes in the future, you must update the information - it won't automatically update.

- Choose **Next page**. 8.
- For **Review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create compute environment**.

# Tutorial: Create an unmanaged compute environment using Amazon EC2 resources

Complete the following steps to create an unmanaged compute environment using Amazon Elastic Compute Cloud (Amazon EC2) resources.

1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.

- From the navigation bar, select the Amazon Web Services Region to use. 2.
- 3. On the Compute Environments page, choose Create.
- 4. Configure the environment.
  - For Compute environment configuration, choose Amazon Elastic Compute Cloud (Amazon EC2).
  - For **Orchestration type**, choose **Unmanaged**.
- For **Name**, specify a unique name for your compute environment. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores ( ).
- 6. For **Service role**, choose a role that lets the Amazon Batch service make calls to the required Amazon API operations on your behalf.



### Note

You can't use AWSServiceRoleForBatch for unmanaged compute environments.

- For **Maximum vCPUs**, choose the maximum number of vCPUs that your compute environment can scale out to, regardless of job queue demand.
- (Optional) Expand Tags. To add a tag, choose Add tag. Then, enter a Key name and optional Value. Choose Add tag. For more information, see Tag your Amazon Batch resources.
- 9. Choose **Next page**.
- 10. For **Review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create compute environment**.

# Tutorial: Create a managed compute environment using Amazon EKS resources

Complete the following steps to create a managed compute environment using Amazon Elastic Kubernetes Service (Amazon EKS) resources.

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- In the navigation pane, choose **Compute environments**. 3.
- 4. Choose Create.

For Compute environment configuration, choose Amazon Elastic Kubernetes Service (Amazon EKS).

- For **Name**, specify a unique name for your compute environment. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- For **Instance role**, choose an existing instance profile that has the required IAM permissions attached.



### Note

To create a compute environment in the Amazon Batch console, choose an instance profile that has the eks:ListClusters and eks:DescribeCluster permissions.

- For **EKS cluster**, choose an existing Amazon EKS cluster. 8.
- For Namespace, enter a Kubernetes namespace to group your Amazon Batch processes in the cluster.
- 10. (Optional) Expand **Tags**. Choose **Add tag** and then enter a key-value pair.
- 11. Choose **Next page**.
- 12. (Optional) For Use EC2 Spot Instances, turn on Enable using Spot instances to use Amazon EC2 Spot Instances.
- 13. (Spot only) For Maximum % on-demand price, choose the maximum percentage that a Spot Instance price can be when compared with the On-Demand price for that instance type before instances are launched. For example, if your maximum price is 20%, then the Spot price must be less than 20% of the current On-Demand price for that EC2 instance. You always pay the lowest (market) price and never more than your maximum percentage. If you leave this field empty, the default value is 100% of the On-Demand price.
- 14. (Spot only) For **Spot fleet role**, choose the Amazon EC2 Spot fleet IAM role for the SPOT compute environment.



### Important

This role is required if the allocation strategy is set to BEST\_FIT or not specified.

15. (Optional) For Minimum vCPUs, choose the minimum number of vCPUs that your compute environment maintains, regardless of job queue demand.

16. (Optional) For Maximum vCPUs, choose the maximum number of vCPUs that your compute environment can scale out to, regardless of job gueue demand.

17. For **Allowed instance types**, choose the Amazon EC2 instance types that can be launched. You can specify instance families to launch any instance type within those families (for example, c5, c5n, or p3). Or, you can specify specific sizes within a family (for example, c5.8xlarge). Metal instance types aren't in the instance families. For example, c5 doesn't include c5.metal. You can also choose optimal to select instance types (from the C4, M4, and R4 instance families) because you need that match the demand of your job queues.

### Note

When you create a compute environment, the instance types that you select for the compute environment must share the same architecture. For example, you can't mix x86 and ARM instances in the same compute environment.

### Note

Amazon Batch scales GPUs based on the required amount in your job queues. To use GPU scheduling, the compute environment must include instance types from the p3, p4, p5, p6, q3, q3s, q4, q5, or q6 families.

### Note

Currently, optimal uses instance types from the C4, M4, and R4 instance families. In Amazon Web Services Regions that don't have instance types from those instance families, instance types from the C5, M5, and R5 instance families are used.

- (Optional) Expand Additional configuration.
  - (Optional) For **Placement group**, enter a placement group name to group resources in the a. compute environment.
  - For Allocation strategy, choose BEST\_FIT\_PROGRESSIVE. b.
  - (Optional) For Amazon Machine Images (AMIs) Configuration, choose Add amazon c. machine images (amis) configuration.

You can use either an Amazon EKS-optimized Amazon Linux AMI or a custom AMI.

- i. To use an Amazon EKS-optimized Amazon Linux AMI:
  - A. For **Image type** choose one of the following:
    - Amazon Linux 2: Default for all Amazon Graviton-based instance families (for example, C6g, M6g, R6g, and T4g) and can be used for all non-GPU instance types.
    - <u>Amazon Linux 2 (accelerated)</u>: Default for all GPU instance families (for example, P4 and G4) and can be used for all non Amazon Graviton-based instance types.
    - Amazon Linux 2023: Amazon Batch supports Amazon Linux 2023 (AL2023).
    - Amazon Linux 2023 (accelerated): GPU instance families and can be used for all non Amazon Graviton-based instance types.
  - B. For **Kubernetes version** enter in a Kubernetes version number.
- ii. To use a custom AMI:
  - A. For **Image type** choose the AMI type that the custom AMI is based off of:
    - <u>Amazon Linux 2</u>: Default for all Amazon Graviton-based instance families (for example, C6g, M6g, R6g, and T4g) and can be used for all non-GPU instance types.
    - Amazon Linux 2 (accelerated): Default for all GPU instance families (for example, P4 and G4) and can be used for all non Amazon Graviton-based instance types.
    - Amazon Linux 2023: Amazon Batch supports AL2023.
    - Amazon Linux 2023 (accelerated): GPU instance families and can be used for all non Amazon Graviton-based instance types.
  - B. For **Image ID override** enter the custom AMI ID.
  - C. For **Kubernetes version** enter in a Kubernetes version number.
- d. (Optional) For Launch template, choose an existing launch template.
- e. (Optional) For Launch template version, enter \$Default, \$Latest, or a version number.
- f. (Optional) For **Override launch template**, to add an override choose **Add override Launch template**:

(Optional) For Launch template choose the launch template to add the override to. i.

- ii. (Optional) For Launch template version choose the version number of the launch template, \$Default, or \$Latest.
- iii. (Optional) For **Target instances types** choose the instance type or family that this override should be applied to. This can target only instance types and families that are included in Allowed instance types.
- iv. (Optional) For userdataType choose the EKS node initialization. Only use this field if you have an AMI specified in either the Launch Template or as a Launch Template Override. Choose **EKS\_NODEADM** for custom AMIs based on EKS\_AL2023 or EKS\_AL2023\_NVIDIA or **EKS\_BOOSTRAP\_SH** for EKS\_AL2 and EKS\_AL\_NVIDIA. The default value is **EKS\_BOOSTRAP\_SH**.

You would use userdataType when you have a mixed environment where you're using both AL2 and AL2023-based custom AMIs in the same compute environment.

- 19. Choose Next page.
- 20. For Virtual Private Cloud (VPC) ID, choose a VPC where to launch the instances.
- 21. For **Subnets**, choose the subnets to use. By default, all subnets within the selected VPC are available.



### Note

Amazon Batch on Amazon EKS supports Local Zones. For more information, see Amazon EKS and Amazon Local Zones in the Amazon EKS User Guide.

- 22. (Optional) For **Security groups**, choose a security group to attach to your instances. By default, the default security group for your VPC is selected.
- 23. Choose **Next page**.
- 24. For **Review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose Create compute environment.

# **Resource: Compute environment template**

The following example shows an empty compute environment template. You can use this template to create your compute environment that can then be saved to a file and used with

the Amazon CLI --cli-input-json option. For more information about these parameters, see CreateComputeEnvironment in the Amazon Batch API Reference.



You can generate a compute environment template with the following Amazon CLI command.

\$ aws batch create-compute-environment --generate-cli-skeleton

```
{
    "computeEnvironmentName": "",
    "type": "UNMANAGED",
    "state": "DISABLED",
    "unmanagedvCpus": 0,
    "computeResources": {
        "type": "EC2",
        "allocationStrategy": "BEST_FIT_PROGRESSIVE",
        "minvCpus": 0,
        "maxvCpus": 0,
        "desiredvCpus": 0,
        "instanceTypes": [
        ],
        "imageId": "",
        "subnets": [
            11.11
        "securityGroupIds": [
        "ec2KeyPair": "",
        "instanceRole": "",
        "tags": {
            "KeyName": ""
        },
        "placementGroup": "",
        "bidPercentage": 0,
        "spotIamFleetRole": "",
        "launchTemplate": {
```

```
"launchTemplateId": "",
             "launchTemplateName": "",
            "version": ""
        },
        "ec2Configuration": [
            {
                 "imageType": "",
                 "imageIdOverride": "",
                 "imageKubernetesVersion": ""
            }
        ]
    },
    "serviceRole": "",
    "tags": {
        "KeyName": ""
    },
    "eksConfiguration": {
        "eksClusterArn": "",
        "kubernetesNamespace": ""
    }
}
```

# **Compute resource AMIs**

By default, Amazon Batch managed compute environments use a recent, approved version of the Amazon ECS optimized AMI for compute resources. However, you might want to create your own AMI to use for your managed and unmanaged compute environments. If you require any of the following, we recommend you create your own AMI:

- Increasing the storage size of your AMI root or data volumes
- Adding instance storage volumes for supported Amazon EC2 instance types
- Customizing the Amazon ECS container agent
- · Customizing Docker
- Configuring a GPU workload AMI to allow containers to access GPU hardware on supported Amazon EC2 instance types

Compute resource AMIs 88



### Note

After a compute environment is created, Amazon Batch doesn't upgrade the AMIs in the compute environment. Amazon Batch also doesn't update the AMIs in your compute environment when a newer version of the Amazon ECS optimized AMI is available. You're responsible for the management of the guest operating system. This includes any updates and security patches. You're also responsible for any additional application software or utilities that you install on the compute resources. To use a new AMI for your Amazon Batch jobs, do the following:

- 1. Create a new compute environment with the new AMI.
- 2. Add the compute environment to an existing job queue.
- 3. Remove the earlier compute environment from your job queue.
- 4. Delete the earlier compute environment.

In April 2022, Amazon Batch added enhanced support for updating compute environments. For more information, see Updating compute environments. To use the enhanced updating of compute environments to update AMIs, follow these rules:

- Either don't set the service role (serviceRole) parameter or set it to the AWSServiceRoleForBatch service-linked role.
- Set the allocation strategy (allocationStrategy) parameter to BEST\_FIT\_PROGRESSIVE, SPOT\_CAPACITY\_OPTIMIZED, or SPOT\_PRICE\_CAPACITY\_OPTIMIZED.
- Set the update to latest image version (updateToLatestImageVersion) parameter to true.
- Don't specify an AMI ID in imageId, imageIdOverride (in ec2Configuration), or in the launch template (launchTemplate). When you don't specify an AMI ID, Amazon Batch selects the latest Amazon ECS optimized AMI that Amazon Batch supports at the time the infrastructure update is initiated. Alternatively, you can specify the AMI ID in the imageId or imageIdOverride parameters. Or, you can specify the launch template that's identified by the LaunchTemplate properties. Changing any of these properties starts an infrastructure update. If the AMI ID is specified in the launch template, the AMI ID can't be replaced by specifying an AMI ID in either the imageId or imageIdOverride parameters. The AMI ID can only be replaced by specifying a different launch template. If the launch template version is set to \$Default or \$Latest, the AMI ID can be replaced

Compute resource AMIs 89

by setting either a new default version for the launch template (if \$Default) or by adding a new version to the launch template (if \$Latest).

If these rules are followed, any update that starts an infrastructure update causes the AMI ID to be re-selected. If the version setting in the launch template (launchTemplate) is set to \$Latest or \$Default, the latest or default version of the launch template is evaluated up at the time of the infrastructure update, even if the launchTemplate wasn't updated.

### **Topics**

- Compute resource AMI specification
- Tutorial: Create a compute resource AMI
- Use a GPU workload AMI
- Amazon Linux deprecation

# **Compute resource AMI specification**

The basic Amazon Batch compute resource AMI specification consists of the following:

### Required

 A modern Linux distribution that's running at least version 3.10 of the Linux kernel on an HVM virtualization type AMI. Windows containers aren't supported.

### Important

Multi-node parallel jobs can only run on compute resources that were launched on an Amazon Linux instance with the ecs-init package installed. We recommend that you use the default Amazon ECS optimized AMI when you create your compute environment. You can do this by not specifying a custom AMI. For more information, see Multi-node parallel jobs.

• The Amazon ECS container agent. We recommend that you use the latest version. For more information, see Installing the Amazon ECS Container Agent in the Amazon Elastic Container Service Developer Guide.

• The awslogs log driver must be specified as an available log driver with the ECS\_AVAILABLE\_LOGGING\_DRIVERS environment variable when the Amazon ECS container agent is started. For more information, see Amazon ECS Container Agent Configuration in the Amazon Elastic Container Service Developer Guide.

 A Docker daemon that's running at least version 1.9, and any Docker runtime dependencies. For more information, see Check runtime dependencies in the Docker documentation.



### Note

We recommend the Docker version that ships with and is tested with the corresponding Amazon ECS agent version that you're using. Amazon ECS provides a changelog for the Linux variant of the Amazon ECS-optimized AMI on GitHub. For more information, see Changelog.

### Recommended

 An initialization and nanny process to run and monitor the Amazon ECS agent. The Amazon ECS optimized AMI uses the ecs-init upstart process, and other operating systems might use systemd. For more information and examples, see Example container instance User Data Configuration Scripts in the Amazon Elastic Container Service Developer Guide. For more information about ecs-init, see the ecs-init project on GitHub. At a minimum, managed compute environments require the Amazon ECS agent to start at boot. If the Amazon ECS agent isn't running on your compute resource, then it can't accept jobs from Amazon Batch.

The Amazon ECS optimized AMI is preconfigured with these requirements and recommendations. We recommend that you use the Amazon ECS optimized AMI or an Amazon Linux AMI with the ecs-init package that's installed for your compute resources. Choose another AMI if your application requires a specific operating system or a Docker version that's not yet available in those AMIs. For more information, see Amazon ECS-Optimized AMI in the Amazon Elastic Container Service Developer Guide.

# **Tutorial: Create a compute resource AMI**

You can create your own custom compute resource AMI to use for your managed and unmanaged compute environments. For instructions, see the Compute resource AMI specification. Then, after

you created a custom AMI, you can create a compute environment that uses that AMI that you can associate a job gueue with. Last, start submitting jobs to that gueue.

### To create a custom compute resource AMI

Choose a base AMI to start from. The base AMI must use HVM virtualization. The base AMI 1. can't be a Windows AMI.



### Note

The AMI that you choose for a compute environment must match the architecture of the instance types that you intend to use for that compute environment. For example, if your compute environment uses A1 instance types, the compute resource AMI that you choose must support ARM instances. Amazon ECS vends both x86 and ARM versions of the Amazon ECS optimized Amazon Linux 2 AMI. For more information, see Amazon ECS optimized Amazon Linux 2 AMI in the Amazon Elastic Container Service Developer Guide.

The Amazon ECS optimized Amazon Linux 2 AMI is the default AMI for compute resources in managed compute environments. The Amazon ECS optimized Amazon Linux 2 AMI is preconfigured and tested on Amazon Batch by Amazon engineers. It's a minimal AMI that you can get started with and to get your compute resources that are running on Amazon guickly. For more information, see Amazon ECS Optimized AMI in the Amazon Elastic Container Service Developer Guide.

Alternatively, you can choose another Amazon Linux 2 variant and install the ecs-init package with the following commands. For more information, see Installing the Amazon ECS container agent on an Amazon Linux 2 EC2 instance in the Amazon Elastic Container Service Developer Guide:

```
$ sudo amazon-linux-extras disable docker
$ sudo amazon-linux-extras install ecs-init
```

For example, if you want to run GPU workloads on your Amazon Batch compute resources, you can start with the Amazon Linux Deep Learning AMI. Then, configure the AMI to run Amazon Batch jobs. For more information, see Use a GPU workload AMI.

### Important

You can choose a base AMI that doesn't support the ecs-init package. However, if you do, you must configure a way to start the Amazon ECS agent at boot and keep it running. You can also view several example user data configuration scripts that use systemd to start and monitor the Amazon ECS container agent. For more information, see Example container instance user data configuration scripts in the Amazon Elastic Container Service Developer Guide.

- 2. Launch an instance from your selected base AMI with the appropriate storage options for your AMI. You can configure the size and number of attached Amazon EBS volumes, or instance storage volumes if the instance type you selected supports them. For more information, see Launching an Instance and Amazon EC2 Instance Store in the Amazon EC2 User Guide.
- Connect to your instance with SSH and perform any necessary configuration tasks. This might include any or all of the following steps:
  - Installing the Amazon ECS container agent. For more information, see Installing the Amazon ECS Container Agent in the Amazon Elastic Container Service Developer Guide.
  - Configuring a script to format instance store volumes.
  - Adding instance store volume or Amazon EFS file systems to the /etc/fstab file so that they're mounted at boot.
  - Configuring Docker options, such as enabling debugging or adjusting base image size.
  - Installing packages or copying files.

For more information, see Connecting to Your Linux Instance Using SSH in the Amazon EC2 User Guide.

- If you started the Amazon ECS container agent on your instance, you must stop it and remove any persistent data checkpoint files before creating your AMI. Otherwise, if you don't do this, the agent doesn't start on instances that are launched from your AMI.
  - Stop the Amazon ECS container agent.
    - Amazon ECS-optimized Amazon Linux 2 AMI:

sudo systemctl stop ecs

• Amazon ECS-optimized Amazon Linux AMI:

```
sudo stop ecs
```

b. Remove the persistent data checkpoint files. By default, these files are located in the / var/lib/ecs/data/ directory. Use the following command to remove these files, if there are any.

```
sudo rm -rf /var/lib/ecs/data/*
```

5. Create a new AMI from your running instance. For more information, see <u>Creating an Amazon EBS Backed Linux AMI</u> in the *Amazon EC2 User Guide* guide.

### To use your new AMI with Amazon Batch

After the new AMI is created, create a compute environment with the new AMI. To do
this, choose the image type and enter the custom AMI ID in the Image ID override box when
you create the Amazon Batch compute environment. For more information, see the section
called "Tutorial: Create a managed compute environment using Amazon EC2 resources".

### Note

The AMI that you choose for a compute environment must match the architecture of the instance types that you intend to use for that compute environment. For example, if your compute environment uses A1 instance types, the compute resource AMI that you choose must support ARM instances. Amazon ECS vends both x86 and ARM versions of the Amazon ECS optimized Amazon Linux 2 AMI. For more information, see <a href="Mailto:Amazon ECS optimized Amazon Linux 2 AMI">Amazon ECS optimized Amazon Linux 2 AMI</a> in the Amazon Elastic Container Service Developer Guide.

2. Create a job queue and associate your new compute environment. For more information, see Create a job queue.



### Note

All compute environments that are associated with a job queue must share the same architecture. Amazon Batch doesn't support mixing compute environment architecture types in a single job queue.

(Optional) Submit a sample job to your new job queue. For more information, see Reference: Job definition examples, Create a single-node job definition, and Tutorial: submit a job.

### Use a GPU workload AMI

To run GPU workloads on your Amazon Batch compute resources, you must use an AMI with GPU support. For more information, see Working with GPUs on Amazon ECS and Amazon ECSoptimized AMIs in Amazon Elastic Container Service Developer Guide.

In managed compute environments, if the compute environment specifies any p3, p4, p5, p6, q3, g3s, g4, g5, or g6 instance types or instance families, then Amazon Batch uses an Amazon ECS GPU optimized AMI.

In unmanaged compute environments, an Amazon ECS GPU-optimized AMI is recommended. You can use the Amazon Command Line Interface or Amazon Systems Manager Parameter Store GetParameter, GetParameters, and GetParametersByPath operations to retrieve the metadata for the recommended Amazon ECS GPU-optimized AMIs.



### Note

The p5 instance family is only supported on versions equal or later than 20230912 of the Amazon ECS GPU-optimized AMI and they are incompatible with p2 and g2 instance types. If you need to use p5 instances, ensure that your compute environment doesn't contain p2 or g2 instances and uses the latest default Batch AMI. Creating a new compute environment will use the latest AMI but If you are updating your compute environment to include p5, you can ensure you are using the latest AMI by setting updateToLatestImageVersion to true in ComputeResource properties. For more information on AMI compatibility with GPU instances, see Working with GPUs on Amazon ECS in Amazon Elastic Container Service Developer Guide.

Use a GPU workload AMI

The following examples show how to use the GetParameter command.

### Amazon CLI

The output includes the AMI information in the Value parameter.

```
{
    "Parameter": {
        "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended",
        "LastModifiedDate": 1555434128.664,
        "Value": "{\"schema_version\":1,\"image_name\":\"amzn2-ami-ecs-gpu-hvm-2.0.20190402-x86_64-ebs\",\"image_id\":\"ami-083c800fe4211192f\",\"os\":\"Amazon Linux 2\",\"ecs_runtime_version\":\"Docker version 18.06.1-ce\",\"ecs_agent_version\":\"1.27.0\"}",
        "Version": 9,
        "Type": "String",
        "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended"
    }
}
```

### Python

```
from __future__ import print_function
import json
import boto3

ssm = boto3.client('ssm', 'us-east-2')

response = ssm.get_parameter(Name='/aws/service/ecs/optimized-ami/amazon-linux-2/
gpu/recommended')
jsonVal = json.loads(response['Parameter']['Value'])
print("image_id = " + jsonVal['image_id'])
print("image_name = " + jsonVal['image_name'])
```

The output only includes the AMI ID and AMI name:

Use a GPU workload AMI 96

```
image_id = ami-083c800fe4211192f
image_name = amzn2-ami-ecs-gpu-hvm-2.0.20190402-x86_64-ebs
```

The following examples demonstrate the use of GetParameters.

#### Amazon CLI

The output includes the full metadata for each of the parameters:

```
{
    "InvalidParameters": [],
    "Parameters": [
        {
            "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
image_id",
            "LastModifiedDate": 1555434128.749,
            "Value": "ami-083c800fe4211192f",
            "Version": 9,
            "Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/gpu/recommended/image_id"
        },
        {
            "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
image_name",
            "LastModifiedDate": 1555434128.712,
            "Value": "amzn2-ami-ecs-gpu-hvm-2.0.20190402-x86_64-ebs",
            "Version": 9,
            "Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/gpu/recommended/image_name"
        }
    ]
}
```

#### Python

The output includes the AMI ID and AMI name, using the full path for the names.

```
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/image_id = ami-083c800fe4211192f /aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/image_name = amzn2-ami-ecs-gpu-hvm-2.0.20190402-x86_64-ebs
```

The following examples show how to use the GetParametersByPath command.

#### Amazon CLI

The output includes the full metadata for all of the parameters under the specified path.

```
"Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/gpu/recommended/ecs_agent_version"
        },
        {
            "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
ecs_runtime_version",
            "LastModifiedDate": 1548368308.213,
            "Value": "Docker version 18.06.1-ce",
            "Version": 1,
            "Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/qpu/recommended/ecs_runtime_version"
        },
        {
            "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
image_id",
            "LastModifiedDate": 1555434128.749,
            "Value": "ami-083c800fe4211192f",
            "Version": 9,
            "Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/gpu/recommended/image_id"
        },
        {
            "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
image_name",
            "LastModifiedDate": 1555434128.712,
            "Value": "amzn2-ami-ecs-gpu-hvm-2.0.20190402-x86_64-ebs",
            "Version": 9,
            "Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/gpu/recommended/image_name"
        },
        {
            "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
os",
            "LastModifiedDate": 1548368308.143,
            "Value": "Amazon Linux 2",
            "Version": 1,
            "Type": "String",
            "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-
ami/amazon-linux-2/gpu/recommended/os"
        },
```

```
{
    "Name": "/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/
schema_version",
    "LastModifiedDate": 1548368307.914,
    "Value": "1",
    "Version": 1,
    "Type": "String",
    "ARN": "arn:aws-cn:ssm:us-east-2::parameter/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/schema_version"
    }
]
```

## Python

```
from __future__ import print_function
import boto3

ssm = boto3.client('ssm', 'us-east-2')

response = ssm.get_parameters_by_path(Path='/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended')
for parameter in response['Parameters']:
    print(parameter['Name'] + " = " + parameter['Value'])
```

The output includes the values of all the parameter names at the specified path, using the full path for the names.

```
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/ecs_agent_version = 1.27.0
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/ecs_runtime_version = Docker version 18.06.1-ce
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/image_id = ami-083c800fe4211192f
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/image_name = amzn2-ami-ecs-gpu-hvm-2.0.20190402-x86_64-ebs
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/os = Amazon Linux 2
/aws/service/ecs/optimized-ami/amazon-linux-2/gpu/recommended/schema_version = 1
```

For more information, see Retrieving Amazon ECS-Optimized AMI Metadata in the Amazon Elastic Container Service Developer Guide.

# **Amazon Linux deprecation**

The Amazon Linux AMI (also called Amazon Linux 1) reached its end of life on December 31, 2023. Amazon Batch has ended support for Amazon Linux AMI as it will not receive any security updates or bug fixes starting January 1, 2024. For more information about the Amazon Linux end-of-life, see AL FAQ.

We recommend that you update existing Amazon Linux based compute environments to Amazon Linux 2023 to prevent unforeseen workload interruptions, and continue to receive security and other updates.

Your compute environments using the Amazon Linux AMI may continue functioning beyond the December 31, 2023 end-of-life date. However, these compute environments will no longer receive any new software updates, security patches, or bug fixes from Amazon. It is your responsibility to maintain these compute environments on the Amazon Linux AMI after end-of-life. We recommend migrating Amazon Batch compute environments to Amazon Linux 2023 or Amazon Linux 2 to maintain optimal performance and security.

For help migrating Amazon Batch from the Amazon Linux AMI to Amazon Linux 2023 or Amazon Linux 2, see Updating compute environments - Amazon Batch.

# Use Amazon EC2 launch templates with Amazon Batch

Amazon Batch supports using Amazon EC2 launch templates with your EC2 compute environments. With launch templates, you can modify the default configuration of your Amazon Batch compute resources without needing to create customized AMIs.



#### Note

Launch templates aren't supported on Amazon Fargate resources.

You must create a launch template before you can associate it with a compute environment. You can create a launch template in the Amazon EC2 console. Or, you can use the Amazon CLI or an Amazon SDK. For example, the following JSON file represents a launch template that resizes the

Amazon Linux deprecation 101

Docker data volume for the default Amazon Batch compute resource AMI and also sets it to be encrypted.

You can create the previous launch template by saving the JSON to a file that's called lt-data. json and running the following Amazon CLI command.

```
aws ec2 --region <region> create-launch-template --cli-input-json file://lt-data.json
```

For more information about launch templates, see <u>Launching an Instance from a Launch Template</u> in the *Amazon EC2 User Guide*.

If you use a launch template to create your compute environment, you can move the following existing compute environment parameters to your launch template:

# Note

Suppose that any of these parameters (except the Amazon EC2 tags) are specified both in the launch template and in the compute environment configuration. Then, the compute environment parameters take precedence. Amazon EC2 tags are merged between the launch template and the compute environment configuration. If there's a collision on the tag's key, the value in the compute environment configuration takes precedence.

Amazon EC2 key pair

- Amazon EC2 AMI ID
- Security group IDs
- Amazon EC2 tags

The following launch template parameters are **ignored** by Amazon Batch:

Instance type (specify your desired instance types when you create your compute environment)

- Instance role (specify your desired instance role when you create your compute environment)
- Network interface subnets (specify your desired subnets when you create your compute environment)
- Instance market options (Amazon Batch must control Spot Instance configuration)
- Disable API termination (Amazon Batch must control instance lifecycle)

Amazon Batch only updates the launch template with a new launch template version during infrastructure updates. For more information, see Updating compute environments.

# Default and override launch templates

You can define a default launch template for the compute environment and an override launch template for specific instance types and families. This can be useful to you so that the default template is used for the majority of instance types in the compute environments.

The substitution variables \$Default and \$Latest can be used instead of naming a specific version. If you do not provide an override launch template, the default launch template is automatically applied.

If you use either the \$Default or \$Latest variable, Amazon Batch will apply the current information at the time that the compute environment is created. If the default or latest version changes in the future, you must update the information through UpdateComputeEnvironment or through the Amazon Web Services Management Console - Amazon Batch.

To provide additional flexibility, you can define override launch templates are applied to specific compute instance types or families.



## Note

You can specify up to ten (10) override launch templates per compute environment.

Use the targetInstanceTypes parameter to select the instance type or family that should use this override launch template. The instance type or family must be first identified by the instanceTypes parameter.

If you define launch template overrides and decide to remove them later, you can pass an empty array to unset the <a href="overrides">overrides</a> parameter in the <a href="UpdateComputeEnvironment">UpdateComputeEnvironment</a> API operation. You can also choose to not include the overrides parameter when submitting the <a href="UpdateComputeEnvironment">UpdateComputeEnvironment</a> API operation. For more information see, <a href="LaunchTemplateSpecification.overrides">LaunchTemplateSpecification.overrides</a>

For more information, see

<u>LaunchTemplateSpecificationOverride.targetInstanceTypes</u> in the Amazon Batch API Reference guide.

# Amazon EC2 user data in launch templates

You can supply Amazon EC2 user data in your launch template that's run by <u>cloud-init</u> when your instances launch. Your user data can perform common configuration scenarios, including but not limited to the following:

- Including users or groups
- Installing packages
- Creating partitions and file systems

Amazon EC2 user data in launch templates must be in the MIME multi-part archive format. This is because your user data is merged with other Amazon Batch user data that's required to configure your compute resources. You can combine multiple user data blocks together into a single MIME multi-part file. For example, you might want to combine a cloud boothook that configures the Docker daemon with a user data shell script that writes configuration information for the Amazon ECS container agent.

If you're using Amazon CloudFormation, the <u>AWS::CloudFormation::Init</u> type can be used with the <u>cfn-init</u> helper script to perform common configuration scenarios.

A MIME multi-part file consists of the following components:

The content type and part boundary declaration: Content-Type: multipart/mixed;
 boundary="==BOUNDARY=="

- The MIME version declaration: MIME-Version: 1.0
- One or more user data blocks that contain the following components:
  - The opening boundary that signals the beginning of a user data block: --==BOUNDARY==. You must keep the line before this boundary blank.
  - The content type declaration for the block: Content-Type: text/cloud-config; charset="us-ascii". For more information about content types, see the <u>Cloud-Init</u> documentation. You must keep the line after the content type declaration blank.
  - The content of the user data, such as a list of shell commands or cloud-init directives.
- The closing boundary that signals the end of the MIME multi-part file: --==BOUNDARY==--. You must keep the line before the closing boundary blank.

# Note

If you add user data to a launch template in the Amazon EC2 console, you can paste it in as plaintext. Or, you can upload it from a file. If you use the Amazon CLI or an Amazon SDK, you must first base64 encode the user data and submit that string as the value of the UserData parameter when you call CreateLaunchTemplate, as shown in this JSON file.

```
{
    "LaunchTemplateName": "base64-user-data",
    "LaunchTemplateData": {
        "UserData":
    "ewogICAgIkxhdW5jaFRlbXBsYXRlTmFtZSI6ICJpbmNyZWFzZS1jb250YWluZXItdm9sdW..."
    }
}
```

## **Topics**

• Reference: Amazon EC2 launch template examples

# Reference: Amazon EC2 launch template examples

The following are example MIME multi-part files that you can use to create your own templates.

# **Examples**

Example: Mount an existing Amazon EFS file system

- Example: Override default Amazon ECS container agent configuration
- Example: Mount an existing Amazon FSx for Lustre file system

# **Example: Mount an existing Amazon EFS file system**

## Example

This example MIME multi-part file configures the compute resource to install the amazon-efsutils package and mount an existing Amazon EFS file system at /mnt/efs.

```
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary="==MYBOUNDARY=="

--==MYBOUNDARY==
Content-Type: text/cloud-config; charset="us-ascii"

packages:
    amazon-efs-utils

runcmd:
    file_system_id_01=fs-abcdef123
    efs_directory=/mnt/efs

- mkdir -p ${efs_directory}
    echo "${file_system_id_01}:/ ${efs_directory} efs tls,_netdev" >> /etc/fstab
    mount -a -t efs defaults

--==MYBOUNDARY==--
```

# **Example: Override default Amazon ECS container agent configuration**

## **Example**

This example MIME multi-part file overrides the default Docker image cleanup settings for a compute resource.

```
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary="==MYBOUNDARY=="

--==MYBOUNDARY==
Content-Type: text/x-shellscript; charset="us-ascii"
```

```
#!/bin/bash
echo ECS_IMAGE_CLEANUP_INTERVAL=60m >> /etc/ecs/ecs.config
echo ECS_IMAGE_MINIMUM_CLEANUP_AGE=60m >> /etc/ecs/ecs.config
--==MYBOUNDARY==--
```

# Example: Mount an existing Amazon FSx for Lustre file system

## **Example**

This example MIME multi-part file configures the compute resource to install the <code>lustre2.10</code> package from the Extras Library and mount an existing FSx for Lustre file system at <code>/scratch</code> and a mount name of <code>fsx</code>. This example is for Amazon Linux 2. For installation instructions for other Linux distributions, see <code>Installing the Lustre Client</code> in the <code>Amazon FSx for Lustre User Guide</code>. For more information, see <code>Mounting your Amazon FSx file system automatically</code> in the <code>Amazon FSx for Lustre User Guide</code>.

```
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary="==MYBOUNDARY=="

--==MYBOUNDARY==
Content-Type: text/cloud-config; charset="us-ascii"

runcmd:
- file_system_id_01=fs-0abcdef1234567890
- region=us-east-2
- fsx_directory=/scratch
- amazon-linux-extras install -y lustre2.10
- mkdir -p ${fsx_directory}
- mount -t lustre ${file_system_id_01}.fsx.${region}.amazonaws.com@tcp:fsx
${fsx_directory}
--==MYBOUNDARY==--
```

In the <u>volumes</u> and <u>mountPoints</u> members of the container properties the mount points must be mapped into the container.

```
{
    "volumes": [
        {
```

# EC2 configurations

Amazon Batch uses Amazon ECS optimized AMIs for EC2 and EC2 Spot compute environments. The default is <a href="Marzon Linux"><u>Amazon Linux 2</u></a> (ECS\_AL2). Before March 31, 2021, this default was <a href="Amazon Linux"><u>Amazon Linux 2</u></a> (ECS\_AL1) for non-GPU, non Amazon Graviton instances.



Amazon Batch also supports Amazon Linux 2023.

The Amazon Linux AMI (also called Amazon Linux 1) reached its end of life on December 31, 2023. Amazon Batch has ended support for Amazon Linux AMI as it will not receive any security updates or bug fixes starting January 1, 2024. For more information about the Amazon Linux end-of-life, see AL FAQ.

We recommend that you update existing Amazon Linux based compute environments to Amazon Linux 2023 to prevent unforeseen workload interruptions, and continue to receive security and other updates.

Your compute environments using the Amazon Linux AMI may continue functioning beyond the December 31, 2023 end-of-life date. However, these compute environments will no longer receive any new software updates, security patches, or bug fixes from Amazon. It is your responsibility to maintain these compute environments on the Amazon Linux AMI after end-of-life. We recommend migrating Amazon Batch compute environments to Amazon Linux 2023 or Amazon Linux 2 to maintain optimal performance and security.

EC2 configurations 108

For help migrating Amazon Batch from the Amazon Linux AMI to Amazon Linux 2023 or Amazon Linux 2, see Updating compute environments - Amazon Batch

# Instance type allocation strategies for Amazon Batch

When a managed compute environment is created, Amazon Batch selects instance types from the <u>instanceTypes</u> specified that best fit the needs of the jobs. The allocation strategy defines behavior when Amazon Batch needs additional capacity. This parameter isn't applicable to jobs that run on Fargate resources. Don't specify this parameter.

## BEST\_FIT (default)

Amazon Batch selects an instance type that best fits the needs of the jobs with a preference for the lowest-cost instance type. If additional instances of the selected instance type aren't available, Amazon Batch waits for the additional instances to be available. If there aren't enough instances available, or if the user is reaching the <a href="Amazon EC2 service quotas">Amazon EC2 service quotas</a>, then additional jobs don't run until currently running jobs are complete. This allocation strategy keeps costs lower but can limit scaling. If you're using Spot Fleets with BEST\_FIT, the Spot Fleet IAM Role must be specified. BEST\_FIT isn't supported when updating compute environments. For more information, see <a href="Updating compute environments">Updating compute environments</a>.

# Note

Amazon Batch manages Amazon resources in your account. Compute environments with the BEST\_FIT allocation strategy originally utilized launch configurations by default. However, the use of launch configurations with new Amazon accounts will be restricted over time. Therefore, beginning in late April 2024, newly-created BEST\_FIT compute environments will default to launch templates. If your service role lacks permissions to manage launch templates, Amazon Batch may continue to utilize launch configurations. Existing compute environments will continue to use launch configurations.

#### BEST\_FIT\_PROGRESSIVE

Amazon Batch selects additional instance types that are large enough to meet the requirements of the jobs in the queue. Instance types with a lower cost for each unit vCPU are preferred. If additional instances of the previously selected instance types aren't available, Amazon Batch selects new instance types.



#### Note

For multi-node parallel jobs Amazon Batch chooses the optimal instance type available. If the instance type becomes unavailable due to insufficient capacity, other instance types within the family are not launched.

## SPOT\_CAPACITY\_OPTIMIZED

Amazon Batch selects one or more instance types that are large enough to meet the requirements of the jobs in the queue. Instance types that are less likely to be interrupted are preferred. This allocation strategy is only available for Spot Instance compute resources.

#### SPOT\_PRICE\_CAPACITY\_OPTIMIZED

The price and capacity optimized allocation strategy looks at both price and capacity to select the Spot Instance pools that are the least likely to be interrupted and have the lowest possible price. This allocation strategy is only available for Spot Instance compute resources.



#### Note

We recommend that you use SPOT\_PRICE\_CAPACITY\_OPTIMIZED rather than SPOT\_CAPACITY\_OPTIMIZED in most instances.

The BEST\_FIT\_PROGRESSIVE and BEST\_FIT strategies use On-Demand or Spot Instances, and the SPOT\_CAPACITY\_OPTIMIZED and SPOT\_PRICE\_CAPACITY\_OPTIMIZED strategies use Spot Instances. However, Amazon Batch might need to exceed maxvCpus to meet your capacity requirements. In this event, Amazon Batch never exceeds maxvCpus by more than a single instance.

# Compute resource memory management

When the Amazon ECS container agent registers a compute resource into a compute environment, the agent must determine how much memory the compute resource has available to reserve for your jobs. Because of platform memory overhead and memory occupied by the system kernel, this number is different than the installed memory amount for Amazon EC2 instances. For example, an m4.large instance has 8 GiB of installed memory. However, this doesn't always translate to exactly 8192 MiB of memory available for jobs when the compute resource registers.

110 Memory management

Suppose that you specify 8192 MiB for the job, and none of your compute resources have 8192 MiB or greater of memory available to satisfy this requirement. Then, the job can't be placed in your compute environment. If you're using a managed compute environment, Amazon Batch must launch a larger instance type to accommodate the request.

The default Amazon Batch compute resource AMI also reserves 32 MiB of memory for the Amazon ECS container agent and other critical system processes. This memory isn't available for job allocation. For more information, see Reserve system memory.

The Amazon ECS container agent uses the Docker ReadMemInfo() function to query the total memory available to the operating system. Linux provides command line utilities to determine the total memory.

## **Example - Determine Linux total memory**

The **free** command returns the total memory that's recognized by the operating system.

```
$ free -b
```

The following is example output for an m4.large instance that's running the Amazon ECS-optimized Amazon Linux AMI.

```
total used free shared buffers cached
Mem: 8373026816 348180480 8024846336 90112 25534464 205418496
-/+ buffers/cache: 117227520 8255799296
```

This instance has 8373026816 bytes of total memory. This means that there's 7985 MiB available for tasks.

#### **Topics**

- · Reserve system memory
- Tutorial: View compute resource memory
- Memory and vCPU considerations for Amazon Batch on Amazon EKS

# **Reserve system memory**

If you occupy all of the memory on a compute resource with your jobs, it's possible that your jobs contend with critical system processes for memory and possibly cause a system failure. The Amazon ECS container agent provides a configuration variable that's called

Reserve system memory 111

ECS\_RESERVED\_MEMORY. You can use this configuration variable to remove a specified number of MiB of memory from the pool that's allocated to your jobs. This effectively reserves that memory for critical system processes.

The default Amazon Batch compute resource AMI reserves 32 MiB of memory for the Amazon ECS container agent and other critical system processes.

# **Tutorial: View compute resource memory**

You can view how much memory a compute resource registers with in the Amazon ECS console or with the <a href="DescribeContainerInstances">DescribeContainerInstances</a> API operation. If you're trying to maximize your resource utilization by providing your jobs as much memory as possible for a particular instance type, you can observe the memory available for that compute resource and then assign your jobs that much memory.

#### To view compute resource memory

- 1. Open the console at https://console.amazonaws.cn/ecs/v2.
- 2. Choose **Clusters**, and then choose the cluster that hosts your compute resources to view.
  - The cluster name for your compute environment begins with your compute environment name.
- 3. Choose Infrastructure.
- 4. Under **Container instances**, choose the container instance.
- 5. The **Resources and networking** section shows the registered and available memory for the compute resource.

The **Registered** memory value is what the compute resource registered with Amazon ECS when it was first launched, and the **Available** memory value is what hasn't already been allocated to jobs.

# Memory and vCPU considerations for Amazon Batch on Amazon EKS

In Amazon Batch on Amazon EKS, you can specify the resources that are made available to a container. For example, you can specify requests or limits values for vCPU and memory resources.

The following are constraints for specifying vCPU resources:

- At least one vCPU requests or limits value must be specified.
- One vCPU unit is equivalent to one physical or virtual core.
- The vCPU value must be entered in whole numbers or in increments of 0.25.
- The smallest valid vCPU value is 0.25.
- If both are specified, the requests value must be less than or equal to the limits value. This way, you can configure both soft and hard vCPU configurations.
- vCPU values can't be specified in milliCPU form. For example, 100m isn't a valid value.
- Amazon Batch uses the requests value for scaling decisions. If a requests value isn't specified, the limits value is copied to the requests value.

The following are constraints for specifying memory resources:

- At least one memory requests or limits value must be specified.
- Memory values must be in mebibytes (MiBs).
- If both are specified, the requests value must be equal to the limits value.
- Amazon Batch uses the requests value for scaling decisions. If a requests value is not specified, the limits value is copied to the requests value.

The following are constraints for specifying GPU resources:

- If both are specified, the requests value must be equal to the limits value.
- Amazon Batch uses the requests value for scaling decisions. If a requests value isn't specified, the limits value is copied to the requests value.

# **Example: job definitions**

The following Amazon Batch on Amazon EKS job definition configures soft vCPU shares. This lets Amazon Batch on Amazon EKS use all of the vCPU capacity for the instance type. However, if there are other jobs running, the job is allocated a maximum of 2 vCPUs. Memory is limited to 2 GB.

```
{
    "jobDefinitionName": "MyJobOnEks_Sleep",
    "type": "container",
    "eksProperties": {
```

```
"podProperties": {
            "containers": [
                {
                     "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
                     "command": ["sleep", "60"],
                     "resources": {
                         "requests": {
                             "cpu": "2",
                             "memory": "2048Mi"
                         }
                     }
                }
            ]
        }
    }
}
```

The following Amazon Batch on Amazon EKS job definition has a request value of 1 and allocates a maximum of 4 vCPUs to the job.

```
{
    "jobDefinitionName": "MyJobOnEks_Sleep",
    "type": "container",
    "eksProperties": {
        "podProperties": {
            "containers": [
                    "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
                    "command": ["sleep", "60"],
                    "resources": {
                         "requests": {
                             "cpu": "1"
                         },
                         "limits": {
                             "cpu": "4",
                             "memory": "2048Mi"
                    }
                }
            ]
        }
    }
}
```

The following Amazon Batch on Amazon EKS job definition sets a vCPU limits value of 1 and a memory limits value of 1 GB.

```
{
    "jobDefinitionName": "MyJobOnEks_Sleep",
    "type": "container",
    "eksProperties": {
        "podProperties": {
             "containers": [
                     "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
                     "command": ["sleep", "60"],
                     "resources": {
                         "limits": {
                             "cpu": "1",
                             "memory": "1024Mi"
                         }
                     }
                }
            ]
        }
    }
}
```

When Amazon Batch translates an Amazon Batch on Amazon EKS job into an Amazon EKS pod, Amazon Batch copies thelimits value to the requests value. This is if a requests value isn't specified. When you submit the preceding example job definition, the pod spec is as follows.

```
apiVersion: v1
kind: Pod
...
spec:
...
containers:
   - command:
        - sleep
        - 60
    image: public.ecr.aws/amazonlinux/amazonlinux:2
    resources:
        limits:
            cpu: 1
            memory: 1024Mi
```

```
requests:
  cpu: 1
  memory: 1024Mi
```

# **Node CPU and memory reservations**

Amazon Batch relies on the default logic of the bootstrap. sh file for vCPU and memory reservations. For more information about the bootstrap.sh file, see bootstrap.sh. When you size your vCPU and memory resources, consider the examples that follow.



## Note

If no instances are running, vCPU and memory reservations can initially affect Amazon Batch scaling logic and decision making. After the instances are running, Amazon Batch adjusts the initial allocations.

# **Example: Node CPU reservation**

The CPU reservation value is calculated in millicores using the total number of vCPUs that are available to the instance.

vCPU number	Percentage reserved
1	6%
2	1%
3-4	0.5%
4 and above	0.25%

Using the preceding values, the following is true:

- The CPU reservation value for a c5.large instance with 2 vCPUs is 70 m. This is calculated in the following way: (1\*60) + (1\*10) = 70 m.
- The CPU reservation value for a c5.24xlarge instance with 96 vCPUs is 310 m. This is calculated in the following way: (1\*60) + (1\*10) + (2\*5) + (92\*2.5) = 310 m.

In this example, there are 1930 (calculated 2000-70) millicore vCPU units available to run jobs on a c5.large instance. Suppose your job requires 2 (2\*1000 m) vCPU units, the job doesn't fit on a single c5.large instance. However, a job that requires 1.75 vCPU units fits.

# **Example: Node memory reservation**

The memory reservation value is calculated in mebibytes using the following:

- The instance capacity in mebibytes. For example, an 8 GB instance is 7,748 MiB.
- The kubeReserved value. The kubeReserved value is the amount of memory to reserve for system daemons. The kubeReserved value is calculated in the following way: ((11 \* maximum number of pods that is supported by the instance type) + 255). For information about the maximum number of pods that's supported by an instance type, see eni-max-pods.txt
- The HardEvictionLimit value. When available memory falls below the HardEvictionLimit value, the instance attempts to evict pods.

The formula to calculate the allocatable memory is as follows: (instance\_capacity\_in\_MiB) - (11 \* (maximum\_number\_of\_pods)) - 255 - (HardEvictionLimit\_value.)).

A c5.large instance supports up to 29 pods. For an 8 GB c5.large instance with a HardEvictionLimit value of 100 MiB, the allocatable memory is 7074 MiB. This is calculated in the following way: (7748 - (11 \* 29) - 255 - 100) = 7074 MiB. In this example, an 8,192 MiB job doesn't fit on this instance even though it's an 8 gibibyte (GiB) instance.

#### **DaemonSets**

When you use DaemonSets, consider the following:

- If no Amazon Batch on Amazon EKS instances are running, DaemonSets can initially affect
  Amazon Batch scaling logic and decision making. Amazon Batch initially allocates 0.5 vCPU units
  and 500 MiB for expected DaemonSets. After the instances are running, Amazon Batch adjusts
  the initial allocations.
- If a DaemonSet defines vCPU or memory limits, Amazon Batch on Amazon EKS jobs have fewer resources. We recommend that you keep the number of DaemonSets that are assigned to Amazon Batch jobs as low as possible.

# **Updating compute environments**

After you create a compute environment that uses EC2 resources, you can update many of the settings of the compute environment directly. However, changing some of the settings requires that Amazon Batch replace the instances in the compute environment.

#### Important

Amazon Batch creates and manages multiple Amazon resources on your behalf and within your account, including Amazon EC2 Launch Templates, Amazon EC2 Auto Scaling Groups, Amazon EC2 Spot Fleets, and Amazon ECS Clusters. These managed resources are configured specifically to ensure optimal Amazon Batch operation. Manually modifying these Batch-managed resources, unless explicitly stated in Amazon Batch documentation, may result in unexpected behavior resulting in INVALID Compute Environment, suboptimal instance scaling behavior, delayed workload processing, or unexpected costs. These manual modifications can not be deterministically supported by the Amazon Batch service. Always use the supported Batch APIs or the Batch console to manage your Compute Environments.

# **Updating Amazon Fargate compute environments**

For compute environments that use Fargate resources, you can update the following.

- securityGroupIds
- subnets
- desiredvCpus
- maxvCpus
- minvCpus

Amazon Batch has two update mechanisms. The first is a scaling update where instances are added or removed from the compute environment. The second is an infrastructure update where the instances in the compute environment are replaced. An infrastructure update takes much longer than a scaling update.

If you update compute environments with Amazon Batch, changing only these settings causes a scaling update: desired vCPUs (desiredvCpus), maximum vCPUs (maxvCpus), minimum vCPUs (minvCpus), service role (serviceRole), and state (state).



## Note

When you update the desiredvCpus setting, the value must be between the minvCpus and maxvCpus values.

Additionally, the updated desiredvCpus value must be greater than or equal to the current desiredvCpus value. For more information, see the section called "Error message when you update the desiredvCpus setting".

If any of the following settings are changed in an UpdateComputeEnvironment API action, Amazon Batch initiates an infrastructure update. An infrastructure update requires that the service role is set to AWSServiceRoleForBatch (the default) and that the allocation strategy is BEST\_FIT\_PROGRESSIVE, SPOT\_CAPACITY\_OPTIMIZED, or SPOT PRICE CAPACITY OPTIMIZED. BEST FIT isn't supported. Except for service role, all of the settings that can be changed for a scaling update can also be changed for an infrastructure update.



#### Note

We recommend that you use SPOT\_PRICE\_CAPACITY\_OPTIMIZED rather than SPOT\_CAPACITY\_OPTIMIZED in most instances.

During an infrastructure update, the status of the compute environment changes to UPDATING. New instances are launched using the updated settings. New jobs are scheduled on the new instances. Jobs that are currently running are dispatched according to the infrastructure update policy. For more information, see UpdateComputeEnvironment and UpdatePolicy in the Amazon Batch API Reference.

In the UpdatePolicy data type, consider the following scenarios:



## Note

In these scenarios, the following is true. When an instance is terminated, running jobs are stopped. By default, these jobs aren't retried. To retry one of these jobs after an instance

is terminated, configure a job retry strategy. For more information, see the section called "Automated job retries" in the Amazon Batch User Guide.

- If the terminateJobsOnUpdate setting is set to true, running jobs are terminated during an infrastructure update. The jobExecutionTimeoutMinutes setting is ignored.
- If the terminateJobsOnUpdate setting is set to false, jobs can run for additional time after the infrastructure update occurs. This additional time is configured in the jobExecutionTimeoutMinutes setting. By default, the jobExecutionTimeoutMinutes setting is 30 minutes.

As capacity becomes available in the compute environment, new instances are launched with the updated settings and jobs are started on the new instances. As all of the jobs complete on instances with the old settings, the old instances are terminated. What capacity becoming available means is that desired number of vCPUs is below the maximum number of vCPUs by at least as many vCPUs as required by the smallest instance type.

### Infrastructure updates

An infrastructure update is required to change some settings for a compute environment. If any of the following settings are changed, an infrastructure update is started:

#### Important

The compute environment must use the AWSServiceRoleForBatch service-linked role to make changes that require an infrastructure update.

If the compute environment uses a service-linked role, it can't be changed to use a regular IAM role. Likewise, if the compute environment has a regular IAM role, it can't be changed to use a service-linked role. Therefore, you can only perform infrastructure updates on compute environments that were created by using a service-linked role.

 Allocation strategy (allocationStrategy, must be either BEST\_FIT\_PROGRESSIVE, SPOT\_CAPACITY\_OPTIMIZED, or SPOT\_PRICE\_CAPACITY\_OPTIMIZED. If the original allocation strategy is BEST\_FIT, infrastructure updates aren't supported.)



## Note

We recommend that you use SPOT\_PRICE\_CAPACITY\_OPTIMIZED rather than SPOT\_CAPACITY\_OPTIMIZED in most instances.

- Bid percentage (bidPercentage)
- EC2 configuration (ec2Configuration)
- Key pair (ec2KeyPair)
- Image ID (imageId)
- Instance role (instanceRole)
- Instance types (instanceTypes)
- Launch template (launchTemplate)
- Placement group (placementGroup)
- Security groups (securityGroupIds)
- VPC subnets (subnets)
- EC2 tags (tags)
- Compute environment type (type, can be one of EC2 or SPOT)
- Whether to update to the latest AMI that's supported by Amazon Batch during an infrastructure update updateToLatestImageVersion

# **Updating the AMI ID**

During an infrastructure update, the compute environment's AMI ID might change, depending on whether AMIs are specified in any of these three settings. AMIs are specified in the imageId (in computeResources), imageIdOverride (in ec2Configuration), or the launch template specified in launchTemplate. Suppose that no AMI IDs are specified in any of those settings and the updateToLatestImageVersion setting is true. Then, the latest Amazon ECS optimized AMI supported by Amazon Batch is used for any infrastructure update.

If an AMI ID is specified in at least one of these settings, the update depends on which setting provided the AMI ID used before the update. When you create a compute environment, the priority for selecting an AMI ID is first the launch template, then the imageId setting, and finally the

Updating the AMI ID 121

imageIdOverride setting. However, if the AMI ID that's used came from the launch template, updating either the imageId or imageIdOverride settings doesn't update the AMI ID. The only way to update an AMI ID selected from the launch template is to update the launch template. If the version parameter of the launch template is \$Default or \$Latest, the default or latest version of the specified launch template is evaluated. If a different AMI ID is selected by the default or the latest version of the launch template is selected, that AMI ID is used in the update.

If the launch template was not used to select the AMI ID, the AMI ID that's specified in the imageId or imageIdOverride parameters is used. If both are specified, the AMI ID specified in the imageIdOverride parameter is used.

Suppose that the compute environment uses an AMI ID specified by the imageId, imageIdOverride, or launchTemplate parameters, and you want to use the latest Amazon ECS optimized AMI supported by Amazon Batch. Then, the update must remove the settings that provided AMI IDs. For imageId, this requires specifying an empty string for that parameter. For imageIdOverride, this requires specifying an empty string for the ec2Configuration parameter.

If the AMI ID came from the launch template, you can change to the latest Amazon ECS optimized AMI that's supported by Amazon Batch by either one of the following ways:

- Remove the launch template by specifying an empty string for the launchTemplateId or launchTemplateName parameter. This removes the entire launch template, rather than the AMI ID alone.
- If the updated version of the launch template doesn't specify an AMI ID, the updateToLatestImageVersion parameter must be set to true.

# **Fargate compute environments**

Fargate is a technology that you can use with Amazon Batch to run <u>containers</u> without having to manage servers or clusters of Amazon EC2 instances. With Fargate, you no longer have to provision, configure, or scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing.

When you run your jobs with Fargate resources, you package your application in containers, specify the CPU and memory requirements, define networking and IAM policies, and launch the application. Each Fargate job has its own isolation boundary and does not share the underlying kernel, CPU resources, memory resources, or elastic network interface with another job.

#### **Topics**

- When to use Fargate
- Job definitions on Fargate
- Job queues on Fargate
- Compute environments on Fargate

# When to use Fargate

We recommend using Fargate in most scenarios. Fargate launches and scales the compute to closely match the resource requirements that you specify for the container. With Fargate, you don't need to over-provision or pay for additional servers. You also don't need to worry about the specifics of infrastructure-related parameters such as instance type. When the compute environment needs to be scaled up, jobs that run on Fargate resources can get started more quickly. Typically, it takes a few minutes to spin up a new Amazon EC2 instance. However, jobs that run on Fargate can be provisioned in about 30 seconds. The exact time required depends on several factors, including container image size and number of jobs.

However, we recommend that you use Amazon EC2 if your jobs require any of the following:

- More than 16 vCPUs
- · More than 120 gibibytes (GiB) of memory
- A GPU
- A custom Amazon Machine Image (AMI)
- Any of the <u>linuxParameters</u> parameters

If you have a large number of jobs, we recommend that you use Amazon EC2 infrastructure. For example, if the number of concurrently running jobs exceeds the Fargate throttling limits. This is because, with EC2, jobs can be dispatched at a higher rate to EC2 resources than to Fargate resources. Moreover, more jobs can run concurrently when you use EC2. For more information, see <u>Fargate service quotas</u> in the *Amazon Elastic Container Service Developer Guide*.

# **Job definitions on Fargate**

Amazon Batch jobs on Amazon Fargate don't support all of the job definition parameters that are available. Some parameters are not supported at all, and others behave differently for Fargate jobs.

When to use Fargate 123

The following list describes job definition parameters that are not valid or otherwise restricted in Fargate jobs.

platformCapabilities

Must be specified as FARGATE.

```
"platformCapabilities": [ "FARGATE" ]
```

type

Must be specified as container.

```
"type": "container"
```

Parameters in containerProperties

executionRoleArn

Must be specified for jobs running on Fargate resources. For more information, see <u>IAM Roles</u> for Tasks in the *Amazon Elastic Container Service Developer Guide*.

```
"executionRoleArn": "arn:aws-cn:iam::123456789012:role/ecsTaskExecutionRole"
```

fargatePlatformConfiguration

(Optional, only for Fargate job definitions). Specifies the Fargate platform version, or LATEST for a recent platform version. Possible values for platformVersion are 1.3.0, 1.4.0, and LATEST (default).

```
"fargatePlatformConfiguration": { "platformVersion": "1.4.0" }
```

instanceType, ulimits

Not applicable for jobs running on Fargate resources.

memory, vcpus

These settings must be specified in resourceRequirements privileged

Either don't specify this parameter, or specify false.

Job definitions on Fargate 124

```
"privileged": false
```

#### resourceRequirements

Both memory and vCPU requirements must be specified using <u>supported values</u>. GPU resources aren't supported for jobs that run on Fargate resources.

If you use GuardDuty Runtime Monitoring, there is a slight memory overhead for the GuardDuty security agent. Therefore the memory limit must include the size of the GuardDuty security agent. For information about the GuardDuty security agent memory limits, see <a href="CPU and memory limits">CPU and memory limits</a> in the GuardDuty User Guide. For information about the best practices, see <a href="How do I remediate out of memory errors on my Fargate tasks after enabling Runtime Monitoring">How do I remediate out of memory errors on my Fargate tasks after enabling Runtime Monitoring</a> in the Amazon ECS Developer Guide.

```
"resourceRequirements": [
    {"type": "MEMORY", "value": "512"},
    {"type": "VCPU", "value": "0.25"}
]
```

#### Parameters in linuxParameters

devices, maxSwap, sharedMemorySize, swappiness, tmpfs

Not applicable for jobs that run on Fargate resources.

## Parameters in logConfiguration

logDriver

Only awslogs and splunk are supported. For more information, see <u>Use the awslogs log</u> driver.

Members in networkConfiguration

```
assignPublicIp
```

If the private subnet doesn't have a NAT gateway attached to send traffic to the Internet, <a href="mailto:assignPublicIp">assignPublicIp</a> must be "ENABLED". For more information, see <a href="mailto:Amazon Batch IAM">Amazon Batch IAM</a> execution role.

Job definitions on Fargate 125

# Job queues on Fargate

Amazon Batch job queues on Amazon Fargate are essentially unchanged. The only restriction is that the compute environments that are listed in computeEnvironmentOrder must all be Fargate compute environments (FARGATE or FARGATE\_SPOT). EC2 and Fargate compute environments can't be mixed.

# **Compute environments on Fargate**

Amazon Batch compute environments on Amazon Fargate don't support all of the compute environment parameters that are available. Some parameters are not supported at all. Others have specific requirements for Fargate.

The following list describes compute environment parameters that aren't valid or otherwise restricted in Fargate jobs.

type

This parameter must be MANAGED.

```
"type": "MANAGED"
```

Parameters in the computeResources object

allocationStrategy, bidPercentage, desiredvCpus, imageId, instanceTypes, ec2Configuration, ec2KeyPair, instanceRole, launchTemplate, minvCpus, placementGroup, spotIamFleetRole

These aren't applicable for Fargate compute environments and can't be provided. subnets

If the subnets listed in this parameter don't have NAT gateways attached, the assignPublicIp parameter in the job definition must be set to ENABLED.

tags

This isn't applicable for Fargate compute environments and can't be provided. To specify tags for Fargate compute environments, use the tags parameter that's not in the computeResources object.

Job gueues on Fargate 126

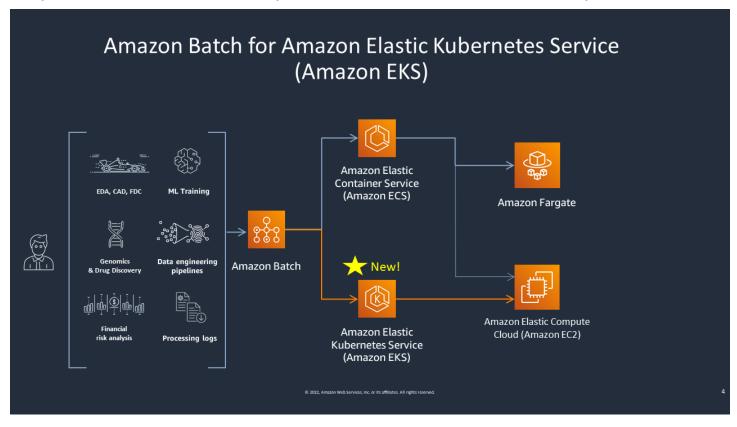
type

This must be either FARGATE or FARGATE\_SPOT.

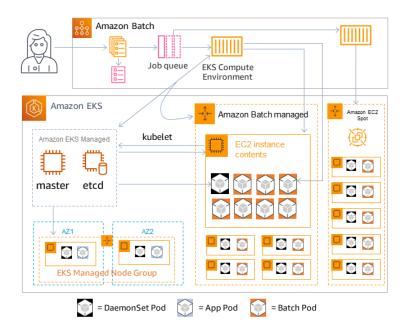
```
"type": "FARGATE_SPOT"
```

# **Amazon EKS compute environments**

<u>Getting started with Amazon Batch on Amazon EKS</u> provides a short guide to creating EKS compute environments. This section provides more details on Amazon EKS compute environments.



Amazon Batch simplifies your batch workloads on Amazon EKS clusters by providing managed batch capabilities. This includes queuing, dependency tracking, managed job retries and priorities, pod management, and node scaling. Amazon Batch can handle multiple Availability Zones and multiple Amazon EC2 instance types and sizes. Amazon Batch integrates several of the Amazon EC2 Spot best practices to run your workloads in a fault-tolerant manner, allowing for fewer interruptions. You can use Amazon Batch to run a handful of overnight jobs or millions of mission-critical jobs with confidence.



Amazon Batch is a managed service that orchestrates batch workloads in your Kubernetes clusters that are managed by Amazon Elastic Kubernetes Service (Amazon EKS). Amazon Batch conducts this orchestration external to your clusters using an "overlay" model. Since Amazon Batch is a managed service, there are no Kubernetes components (for example, Operators or Custom Resources) to install or manage in your cluster. Amazon Batch only needs your cluster to be configured with Role-Based Access Controls (RBAC) that allow Amazon Batch to communicate with the Kubernetes API server. Amazon Batch calls Kubernetes APIs to create, monitor, and delete Kubernetes pods and nodes.

Amazon Batch has built-in scaling logic to scale Kubernetes nodes based on job queue load with optimizations in terms of job capacity allocations. When the job queue is empty, Amazon Batch scales down the nodes to the minimum capacity that you set, which by default is zero. Amazon Batch manages the full lifecycle of these nodes, and decorates the nodes with labels and taints. This way, other Kubernetes workloads aren't placed on the nodes managed by Amazon Batch. The exception to this are DaemonSets, which can target Amazon Batch nodes to provide monitoring and other functionality required for proper execution of the jobs. Additionally, Amazon Batch doesn't run jobs, specifically pods, on nodes in your cluster that it doesn't manage. This way, you can use separate scaling logic and services for other applications on the cluster.

To submit jobs to Amazon Batch, you interact directly with the Amazon Batch API. Amazon Batch translates jobs into podspecs and then creates the requests to place pods on nodes managed by Amazon Batch in your Amazon EKS cluster. You can use tools such as kubectl to view running pods and nodes. When a pod has completed its execution, Amazon Batch deletes the pod it created to maintain a lower load on the Kubernetes system.

You can get started by connecting a valid Amazon EKS cluster with Amazon Batch. Then attach an Amazon Batch job queue to it, and register an Amazon EKS job definition using podspec equivalent attributes. Last, submit jobs using the <a href="SubmitJob">SubmitJob</a> API operation referencing to the job definition. For more information, see <a href="Getting started with Amazon Batch on Amazon EKS">Getting started with Amazon Batch on Amazon EKS</a>.

## **Amazon EKS**

## **Topics**

- Amazon Batch default AMI
- Mixed AMI environments
- Supported Kubernetes versions
- Update the Kubernetes version of the compute environment
- · Shared responsibility of the Kubernetes nodes
- Run a DaemonSet on Amazon Batch managed nodes
- Customize Amazon EKS launch templates
- How to upgrade from EKS AL2 to EKS AL2023

# Amazon Batch default AMI

When you create an Amazon EKS compute environment, you don't need to specify an Amazon Machine Image (AMI). Amazon Batch selects an Amazon EKS optimized AMI based on the Kubernetes version and instance types that are specified in your <a href="CreateComputeEnvironment">CreateComputeEnvironment</a> request. In general, we recommend that you use the default AMI selection. For more information about Amazon EKS optimized AMIs, see<a href="Amazon EKS optimized Amazon Linux AMIs">Amazon Linux AMIs</a> in the Amazon EKS User Guide.

Run the following command to see which AMI type Amazon Batch is selected for your Amazon EKS compute environment. This following example is a non-GPU instance type.

Amazon EKS 129

This following example is a GPU instance type.

# **Mixed AMI environments**

You can use launch template overrides to create compute environments with both Amazon Linux 2 (AL2) and Amazon Linux 2023 (AL2023) AMIs. This is useful for using different AMIs for different architectures

Example configuration using both AMI types:

```
{
  "computeResources": {
    "launchTemplate": {
      "launchTemplateId": "TemplateId",
      "version": "1",
      "userdataType": "EKS_BOOTSTRAP_SH",
      "overrides": [
      {
            "instanceType": "c5.large",
            "imageId": "ami-al2-custom",
            "userdataType": "EKS_BOOTSTRAP_SH"
```

Mixed AMI environments 130

```
},
    {
        "instanceType": "c6a.large",
        "imageId": "ami-al2023-custom",
        "userdataType": "EKS_NODEADM"
     }
    ]
    },
    "instanceTypes": ["c5.large", "c6a.large"]
}
```

# **Supported Kubernetes versions**

Amazon Batch on Amazon EKS currently supports the following Kubernetes versions:

- 1.32
- 1.31
- 1.30
- 1.29
- 1.28
- 1.27
- 1.26
- 1.25

You might see an error message that resembles the following when you use the CreateComputeEnvironment API operation or UpdateComputeEnvironmentAPI operation to create or update a compute environment. This issue occurs if you specify an unsupported Kubernetes version in EC2Configuration.

```
At least one imageKubernetesVersion in EC2Configuration is not supported.
```

To resolve this issue, delete the compute environment and then re-create it with a supported Kubernetes version.

You can perform a minor version upgrade on your Amazon EKS cluster. For example, you can upgrade the cluster from 1.xx to 1.yy even if the minor version isn't supported.

However, the compute environment status might change to INVALID after a major version update. For example, if you perform a major version upgrade from 1.xx to 2.yy. If the major version isn't supported by Amazon Batch, you see an error message that resembles the following.

```
reason=CLIENT_ERROR - ... EKS Cluster version [2.yy] is unsupported
```

# Update the Kubernetes version of the compute environment

With Amazon Batch, you can update the Kubernetes version of a compute environment to support Amazon EKS cluster upgrades. The Kubernetes version of a compute environment is the Amazon EKS AMI version for the Kubernetes nodes that Amazon Batch launches to run jobs. You can perform a Kubernetes version upgrade on their Amazon EKS nodes before or after you update the version of Amazon EKS cluster's control-plane. We recommend that you update the nodes after upgrading the control plane. For more information, see <a href="Updating an Amazon EKS cluster">Updating an Amazon EKS cluster</a> Kubernetes version in the Amazon EKS User Guide.

To upgrade the Kubernetes version of a compute environment, use the UpdateComputeEnvironment API operation.

```
$ aws batch update-compute-environment \
    --compute-environment <compute-environment-name> \
    --compute-resources \
    'ec2Configuration=[{imageType=EKS_AL2,imageKubernetesVersion=1.32}]'
```

# **Shared responsibility of the Kubernetes nodes**

Maintenance of the compute environments is a shared responsibility.

- Don't change or remove Amazon Batch nodes, labels, taints, namespaces, launch templates, or auto scaling groups. Don't add taints to Amazon Batch managed nodes. If you make any of these changes, your compute environment cannot be supported and failures including idle instances occur.
- Don't target your pods to Amazon Batch managed nodes. If you target your pods to the
  managed nodes, broken scaling and stuck job queues occur. Run workloads that don't use
  Amazon Batch on self-managed nodes or managed node groups. For more information, see
  Managed node groups in the Amazon EKS User Guide.
- You can target a DaemonSet to run on Amazon Batch managed nodes. For more information, see Run a DaemonSet on Amazon Batch managed nodes.

Amazon Batch doesn't automatically update compute environment AMIs. It's your responsibility to update them. Run the following command to update your AMIs to the latest AMI version.

```
$ aws batch update-compute-environment \
    --compute-environment < compute-environment-name > \
    --compute-resources 'updateToLatestImageVersion=true'
```

Amazon Batch doesn't automatically upgrade the Kubernetes version. Run the following command to update the Kubernetes version of your computer environment to 1.32.

```
$ aws batch update-compute-environment \
    --compute-environment <compute-environment-name> \
    --compute-resources \
    'ec2Configuration=[{imageType=EKS_AL2,imageKubernetesVersion=1.32}]'
```

When updating to a more recent AMI or the Kubernetes version, you can specify whether to terminate jobs when they're updated (terminateJobsOnUpdate) and how long to wait for before an instance is replaced if running jobs don't finish (jobExecutionTimeoutMinutes.) For more information, see <a href="Updating compute environments">Updating compute environments</a> and the infrastructure update policy (UpdatePolicy) set in the UpdateComputeEnvironment API operation.

# Run a DaemonSet on Amazon Batch managed nodes

Amazon Batch sets taints on Amazon Batch managed Kubernetes nodes. You can target a DaemonSet to run on Amazon Batch managed nodes with the following tolerations.

```
tolerations:
   - key: "batch.amazonaws.com/batch-node"
   operator: "Exists"
```

Another way to do this is with the following tolerations.

```
tolerations:
    - key: "batch.amazonaws.com/batch-node"
    operator: "Exists"
    effect: "NoSchedule"
    - key: "batch.amazonaws.com/batch-node"
    operator: "Exists"
    effect: "NoExecute"
```

# **Customize Amazon EKS launch templates**

Amazon Batch on Amazon EKS supports launch templates. There are constraints on what your launch template can do.

### ▲ Important

- For EKS AL2 AMIs, Amazon Batch runs /etc/eks/bootstrap.sh. Don't run /etc/eks/bootstrap.sh in your launch template or cloud-init user-data scripts. You can add additional parameters besides the --kubelet-extra-args parameter to bootstrap.sh. To do this, set the AWS\_BATCH\_KUBELET\_EXTRA\_ARGS variable in the /etc/aws-batch/batch.config file. See the following example for details.
- For EKS AL2023, Amazon Batch utilizes the <u>NodeConfigSpec</u> from EKS to make instances join the EKS cluster. Amazon Batch populates <u>ClusterDetails</u> in <u>NodeConfigSpec</u> for the EKS cluster and you don't need to specify them.

### Note

We recommend that you do not set any of the follow <u>NodeConfigSpec</u> settings in the launch template as Amazon Batch will override your values. For more information, see <u>Shared responsibility of the Kubernetes nodes</u>.

- Taints
- Cluster Name
- apiServerEndpoint
- certificatAuthority
- CIDR
- Do not create a labels with the prefix batch.amazonaws.com/



### Note

If the launch template is changed after CreateComputeEnvironment is called, UpdateComputeEnvironment must be called to evaluate the version of the launch template for replacement.

### **Topics**

- Add kubelet extra arguments
- Configure the container runtime
- Mount an Amazon EFS volume
- IPv6 support

### Add kubelet extra arguments

Amazon Batch supports adding extra arguments to the kubelet command. For the list of supported parameters, see kubelet in the Kubernetes documentation. In the following example for EKS AL2 AMIs, --node-labels mylabel=helloworld is added to the kubelet command line.

```
MIME-Version: 1.0
      Content-Type: multipart/mixed; boundary="==MYBOUNDARY=="
      --==MYBOUNDARY==
      Content-Type: text/x-shellscript; charset="us-ascii"
      #!/bin/bash
      mkdir -p /etc/aws-batch
      echo AWS_BATCH_KUBELET_EXTRA_ARGS=\"--node-labels mylabel=helloworld\" >> /etc/
aws-batch/batch.config
      --==MYBOUNDARY==--
```

For EKS AL2023 AMIs the file format is YAML. For the list of supported parameters, see NodeConfigSpec in the Kubernetes documentation. In the following example for EKS AL2023 AMIs, --node-labels mylabel=helloworld is added to the kubelet command line.

```
MIME-Version: 1.0
```

```
Content-Type: multipart/mixed; boundary="==MYBOUNDARY=="

--==MYBOUNDARY==
Content-Type: application/node.eks.aws

apiVersion: node.eks.aws/v1alpha1
kind: NodeConfig
spec:
    kubelet:
        flags:
        - --node-labels=mylabel=helloworld

--==MYBOUNDARY==--
```

### Configure the container runtime

You can use the Amazon Batch CONTAINER\_RUNTIME environment variable to configure the container runtime on a managed node. The following example sets the container runtime to containerd when bootstrap.sh runs. For more information, see <a href="mailto:containerd">containerd</a> in the Kubernetes documentation.

If you are using an optimized EKS\_AL2023 or EKS\_AL2023\_NVIDIA AMI you do not need to specify the container runtime as only **containerd** is supported.

### Note

The CONTAINER\_RUNTIME environment variable is equivalent to the --containerruntime option of bootstrap.sh. For more information, see <u>Options</u> in the *Kubernetes* documentation.

```
--==MYBOUNDARY==--
```

#### Mount an Amazon EFS volume

You can use launch templates to mount volumes to the node. In the following example, the cloud-config packages and runcmd settings are used. For more information, see <u>Cloud config</u> examples in the *cloud-init documentation*.

```
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary="==MYBOUNDARY=="

--==MYBOUNDARY==
Content-Type: text/cloud-config; charset="us-ascii"

packages:
    amazon-efs-utils

runcmd:
    file_system_id_01=fs-abcdef123
    efs_directory=/mnt/efs

- mkdir -p ${efs_directory}
    echo "${file_system_id_01}:/ ${efs_directory} efs _netdev,noresvport,tls,iam 0 0"
    >> /etc/fstab
    mount -t efs -o tls ${file_system_id_01}:/ ${efs_directory}
    --==MYBOUNDARY==--
```

To use this volume in the job, it must be added in the <u>eksProperties</u> parameter to RegisterJobDefinition. The following example is a large portion of the job definition.

```
"limits": {
                               "cpu": "1",
                               "memory": "1024Mi"
                          }
                     },
                      "volumeMounts": [
                          {
                               "name": "efs-volume",
                               "mountPath": "/efs"
                          }
                     ]
                 }
             ],
             "volumes": [
                 {
                     "name": "efs-volume",
                     "hostPath": {
                          "path": "/mnt/efs"
                     }
                 }
             ]
        }
    }
}
```

In the node, the Amazon EFS volume is mounted in the /mnt/efs directory. In the container for the Amazon EKS job, the volume is mounted in the /efs directory.

### **IPv6** support

Amazon Batch supports Amazon EKS clusters that have IPv6 addresses. No customizations are required for Amazon Batch support. However, before you begin, we recommend that you review the considerations and conditions that are outlined in <u>Assigning IPv6 addresses to pods and services</u> in the *Amazon EKS User Guide*.

# How to upgrade from EKS AL2 to EKS AL2023

The Amazon EKS optimized AMIs are available in two families based on Amazon Linux 2 (AL2) and Amazon Linux 2023 (AL2023). AL2023 is a new Linux-based operating system designed to provide a secure, stable, and high-performance environment for your cloud applications. For more information about the differences between AL2 and AL2023 see <a href="Upgrade from Amazon Linux 2 to Amazon Linux 2023">Upgrade from Amazon Linux 2 to Amazon Linux 2023</a> in the Amazon EKS User Guide.

Depending on how your compute environment is configured you can use one of the following upgrade path from AL2 to AL2023.

#### Upgrade using Ec2Configuration.ImageType

- If you are not using a launch template or launch template overrides then change <u>Ec2Configuration.ImageType</u> to EKS\_AL2023 or EKS\_AL2023\_NVIDIA and the run UpdateComputeEnvironment.
- If you specify an <a href="Ec2Configuration.lmageIdOverride">Ec2Configuration.lmageType</a> must match the AMI type specified in <a href="Ec2Configuration.lmageIdOverride">Ec2Configuration.lmageIdOverride</a>.

If you mismatch ImageIdOverride and ImageType then the node won't join the cluster.

### **Upgrade using launch templates**

• If you have any kubelet extra arguments defined in a launch template or launch template override, they need to updated to the new <u>kubelet extra arguments format</u>.

If you mismatch the kubelet extra arguments format then the extra arguments aren't applied.

• For AL2023 AMIs, **containerd** is the only supported container runtime. You do not need to specify container runtime for EKS\_AL2023 in the launch template.

You can't specify a customised container runtime with EKS\_AL2023.

• If you use a launch template or launch template override that specifies an AMI based on EKS\_AL2023 then you need to set userdataType to EKS\_NODEADM.

If you mismatch the userdataType and AMI then the node won't join the EKS cluster.

# **Job queues**

Jobs are submitted to a job queue where they reside until they can be scheduled to run in a compute environment. An Amazon account can have multiple job queues. For example, you can create a queue that uses Amazon EC2 On-Demand instances for high priority jobs and another queue that uses Amazon EC2 Spot Instances for low-priority jobs. Job queues have a priority that's used by the scheduler to determine which jobs in which queue should be evaluated for execution first.

### **Topics**

- Create a job queue
- View job queue status
- Fair-share scheduling policies
- Resource-aware scheduling

# Create a job queue

Before you can submit jobs in Amazon Batch, you must create a job queue. When you create a job queue, you associate one or more compute environments to the queue and assign an order of preference.

You also set priority to the job queue that determines the order that the Amazon Batch scheduler places jobs. This means that, if a compute environment is associated with more than one job queue, the job queue with a higher priority is given preference.

#### **Topics**

- Tutorial: Create an Amazon EC2 job queue
- Tutorial: Create a Fargate job queue
- Tutorial: Create an Amazon EKS job queue
- Reference: Job queue template

### **Tutorial: Create an Amazon EC2 job queue**

Complete the following steps to create a job queue for Amazon Elastic Compute Cloud (Amazon EC2).

Create a job queue 140

#### To create an Amazon EC2 job queue

- Open the Amazon Batch console at https://console.amazonaws.cn/batch/. 1.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Job gueues**.
- Choose Create. 4.
- For Orchestration type, choose Amazon Elastic Compute Cloud (Amazon EC2). 5.
- For **Name**, enter a unique name for your job queue. The name can be up to 128 characters 6. long, and can contain uppercase and lowercase letters, numbers, and underscores (\_).
- For **Priority**, enter an whole number value for the job queue's priority. Job queues with a higher priority are run before lower priority job queues that are associated with the same compute environment. Priority is determined in descending order. For example, a job queue with a priority value of 10 is given scheduling preference over a job queue with a priority value of 1.
- (Optional) For Scheduling policy Amazon Resource Name (ARN), choose an existing scheduling policy.
- 9. For **Connected compute environments**, select one or more compute environments from the list to associate with the job queue. Select compute environments in the order that you want the queue to attempt job queue placement. The job scheduler uses the order that you select compute environments in to determine which compute environment starts a given job. Before you can associate them with a job queue, compute environments must be in the VALID state. You can associate up to three compute environments with a job queue. If you don't have an existing compute environment, choose Create compute environment



### Note

All compute environments that are associated with a job queue must share the same provisioning model. Amazon Batch doesn't support mixing provisioning models in a single job queue.

- 10. For **Compute environment order**, choose the up and down arrows to configure order that you want.
- 11. Choose **Create job queue** to finish and create your job queue.

# Tutorial: Create a Fargate job queue

Complete the following steps to create a job queue for Amazon Fargate.

### To create a Fargate job queue

- Open the Amazon Batch console at https://console.amazonaws.cn/batch/. 1.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Job queues**.
- 4. Choose Create.
- 5. For **Orchestration type**, choose **Fargate**.
- For **Name**, enter a unique name for your job queue. The name can be up to 128 characters 6. long, and can contain uppercase and lowercase letters, numbers, and underscores (\_).
- For **Priority**, enter an whole number value for the job queue's priority. Job queues with a higher priority are run before lower priority job queues that are associated with the same compute environment. Priority is determined in descending order. For example, a job queue with a priority value of 10 is given scheduling preference over a job queue with a priority value of 1.
- (Optional) For **Scheduling policy Amazon Resource Name (ARN)**, choose an existing scheduling policy.
- For **Connected compute environments**, select one or more compute environments from the list to associate with the job queue. Select compute environments in the order that you want the queue to attempt job queue placement. The job scheduler uses the order that you select compute environments in to determine which compute environment starts a given job. Before you can associate them with a job queue, compute environments must be in the VALID state. You can associate up to three compute environments with a job queue.

### Note

All compute environments that are associated with a job queue must share the same provisioning model. Amazon Batch doesn't support mixing provisioning models in a single job queue.

- 10. For Compute environment order, choose the up and down arrows to configure order that you want.
- 11. Choose **Create job queue** to finish and create your job queue.

# Tutorial: Create an Amazon EKS job queue

Complete the following steps to create a job queue for Amazon Elastic Kubernetes Service (Amazon EKS).

#### To create an Amazon EKS job queue

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- In the navigation pane, choose Job queues. 3.
- Choose Create. 4.
- 5. For Orchestration type, choose Amazon Elastic Kubernetes Service (Amazon EKS).
- For **Name**, enter a unique name for your job queue. The name can be up to 128 characters 6. long, and can contain uppercase and lowercase letters, numbers, and underscores (\_).
- For **Priority**, enter an integer value for the job queue's priority. Job queues with a higher priority are run before lower priority job queues that are associated with the same compute environment. Priority is determined in descending order. For example, a job queue with a priority value of 10 is given scheduling preference over a job queue with a priority value of 1.
- (Optional) For Scheduling policy Amazon Resource Name (ARN), choose an existing scheduling policy.
- For **Connected compute environments**, select one or more compute environments from the list to associate with the job queue. Select compute environments in the order that you want the queue to attempt job queue placement. The job scheduler uses the order that you select compute environments in to determine which compute environment starts a given job. Before you can associate them with a job queue, compute environments must be in the VALID state. You can associate up to three compute environments with a job queue.



#### Note

All compute environments that are associated with a job queue must share the same provisioning model. Amazon Batch doesn't support mixing provisioning models in a single job queue.



### Note

All compute environments that are associated with a job queue must share the same architecture. Amazon Batch doesn't support mixing compute environment architecture types in a single job queue.

- 10. For **Compute environment order**, choose the up and down arrows to configure order that you want.
- 11. Choose **Create job queue** to finish and create your job queue.

# Reference: Job queue template

The following is an empty job queue template. You can use this template to create your job queue. You can then save this job queue to a file and use it with the Amazon CLI --cli-input-json option. For more information about these parameters, see CreateJobQueue in the Amazon Batch API Reference.

### Note

You can generate a job queue template with the following Amazon CLI command.

```
$ aws batch create-job-queue --generate-cli-skeleton
```

```
{
   "computeEnvironmentOrder": [
         "computeEnvironment": "",
         "order": 0
      }
   ],
   "jobQueueName": "",
   "jobStateTimeLimitActions": [
         "state": "RUNNABLE",
         "action": "CANCEL",
   "maxTimeSeconds": 0,
```

```
"reason": ""

}
],
"priority": 0,
"schedulingPolicyArn": "",
"state": "ENABLED",
"tags": {
    "KeyName": ""
}
```

# View job queue status

After you create a job queue and submit the jobs, it is important to be able to monitor its progress. You can use the **Job details** page to review, manage, and monitor your job queue.

## View job queue information

From the Amazon Batch console, select **Job queues** in navigation pane and choose your desired job queue to view its details. On this page, you can review and manage your job queue and see additional information about the queue's operations, such as the job queue snapshot, job state limits, environment order, tags, and the job queue's JSON code.

# Job queue details

This section provides an overview and maintenance options for the job queue. It is important to note that you can find the Amazon Resource Name (ARN) in this section.

To find this information through the Amazon Command Line Interface, use the DescribeJobQueues operation along with the job queue name, or the corresponding ARN.

## Job queue snapshot

This section provides a static list of the first 100 RUNNABLEjobs that are in queue. You can use the search field to narrow the list by searching for information from any column in the results section. The jobs in the snapshot results area are sorted according to the job queue's run strategy. For first-in-first-out (FIFO) job queues, the ordering of the jobs is based on the submission time. For <u>fair-share scheduling</u> job queues, the ordering of the jobs is based on the job priority and share usage.

View job queue status 145

Because the results are a snapshot of the job queue, the results list doesn't automatically update. To update the list, choose the refresh at the top of the section. Choose the job's name hyperlink to navigate to **Job details** and view the job's status and other related information.

To find this information through the Amazon CLI, use the <u>GetJobQueueSnapshot</u> operation along with the job queue name or the corresponding ARN.

#### Job state limits

Use this tab to review configuration information about the amount of time that a job can remain in a RUNNABLE state before it's canceled.

To find this information through the Amazon CLI, use the <a href="DescribeJobQueues">DescribeJobQueues</a> operation along with the job queue name or the corresponding ARN.

#### **Environment order**

If your job queue runs in multiple environments, this tab provides their order and an overview.

To find this information through the Amazon CLI, use the <u>DescribeJobQueues</u> operation along with the job queue name or the corresponding ARN.

### **Tags**

Use this tab to review and manage tags that are associated to this job queue.

#### **JSON**

Use this tab to copy the JSON code that's associated with this job queue. You can then reuse the JSON for Amazon CloudFormation templates and Amazon CLI scripts.

# Fair-share scheduling policies

The Amazon Batch scheduler evaluates when, where, and how to run jobs that are submitted to a job queue. If you don't specify a scheduling policy when you create a job queue, the Amazon Batch job scheduler defaults to a first-in, first-out (FIFO) strategy. A FIFO strategy might cause important jobs to get "stuck" behind jobs that were submitted earlier. By specifying a different scheduling policy, you can allocate compute resources according to your specific needs.

Fair-share scheduling policies 146



### Note

If you want to schedule the specific order that jobs are run in, use the depends 0n parameter in SubmitJob to specify the dependencies for each job.

If you create a scheduling policy and attach it to a job queue, fair-share scheduling is turned on. If the job queue has a scheduling policy, the scheduling policy determines the order that jobs are run in. For more information, see Use fair-share scheduling policies to assign share identifiers.

### **Topics**

- Use share identifiers to identify workloads
- Use fair-share scheduling policies to assign share identifiers
- Use fair-share scheduling to help schedule jobs
- Tutorial: Create a scheduling policy
- Reference: Scheduling policy template

# Use share identifiers to identify workloads

You can use share identifiers to tag jobs and differentiate between users and workloads. The Amazon Batch scheduler tracks usage for each share identifier by using the (T \* weightFactor)formula, where T is the vCPU usage over time. The scheduler picks jobs with the lowest usage from the share identifier. You can use a share identifier without overriding it.



### Note

Share identifiers are unique within a job queue and are not aggregated across job queues.

You can set fair-share scheduling priority to configure the order that jobs are run in on a share identifier. Jobs with a higher scheduling priority are scheduled first. If you don't specify a fair-share scheduling policy, all jobs that are submitted to the job queue are scheduled in FIFO order. When you submit a job, you can't specify a share identifier or fair-share scheduling priority.

Use share identifiers 147



#### (i) Note

Attached compute resources are allocated equally among all share identifiers unless explicitly overridden.

# Use fair-share scheduling policies to assign share identifiers

You can use scheduling policies to configure how compute resources in a job queue are allocated between users or workloads. Using fair-share scheduling policies, you can assign different share identifiers to workloads or users. Amazon Batch assigns each share identifier a percentage of the total resources that are available during a period of time.

The fair-share percentage is calculated using the shareDecaySeconds and shareDistribution values. You can add time to the fair-share analysis by assigning a share decay time to the policy. Adding time gives more weight to time and less to the defined weight. You can hold compute resources in reserve for share identifiers that aren't active by specifying a compute reservation. For more information, see SchedulingPolicyDetail.

# Use fair-share scheduling to help schedule jobs

Fair-share scheduling provides a set of controls to help schedule jobs.



### Note

For more information about scheduling policy parameters, see SchedulingPolicyDetail.

• Share decay seconds – The period of time (in seconds) that the Amazon Batch scheduler uses to calculate a fair-share percentage for each share identifier. A value of zero indicates that only current usage is measured. A longer decay time gives more weight to time.



#### Note

The period of time for decay is calculated as: shareDecaySeconds + OrderMinutes where OrderMinutes is the time in the order in minutes.

• Compute reservation – Prevents jobs in a single share identifier from using up all the resources that are attached to the job queue. The reserved ratio is

Use scheduling policies 148

(computeReservation/100)^ActiveFairShares where ActiveFairShares is the number of active share identifiers.



#### Note

If a share identifier has jobs in a SUBMITTED, PENDING, RUNNABLE, STARTING, or RUNNING state, it's considered an active share identifier. After the period of time for decay expires, a share identifier is considered inactive.

• Weight factor – The weight factor for the share identifier. The default value is 1. A lower value lets jobs from the share identifier run or gives additional runtime to the share identifier. For example, jobs that use a share identifier with a weight factor of 0.125 (1/8) are assigned eight times the compute resources of jobs that use a share identifier with a weight factor of 1.



#### Note

You only need to define this attribute when you need to update the default weight factor of 1.

When the job queue is active and processing jobs, you can review a list of the first 100 RUNNABLE jobs through the Job queue snapshot. For more information, see View job queue status.

## **Tutorial: Create a scheduling policy**

Before you can create a job queue with a scheduling policy, you must create a scheduling policy. When you create a fair-share scheduling policy, you associate one or more share identifiers or share identifier prefixes with weights for the queue and optionally assign a decay period and compute reservation to the policy.

### To create a scheduling policy

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Region to use.
- 3. In the navigation pane, choose **Scheduling policies**, **Create**.
- For Name, enter a unique name for your scheduling policy. Up to 128 letters (uppercase and 4. lowercase), numbers, hyphens, and underscores are allowed.

5. (Optional) For **Share decay seconds**, enter an integer value for the fair-share scheduling policy's share decay time. A longer share decay time will use considerably more compute resource usage over a longer time when scheduling jobs. This can allow jobs using a share identifier to temporarily use more compute resources than the weight for that share identifier would allow if that share identifier had not recently been using compute resources.

6. (Optional) For **Compute reservation**, enter an integer value for the fair-share scheduling policy's compute reservation. The compute reservation will hold some vCPUs in reserve to be used for share identifiers that are not currently active.

The reserved ratio is (computeReservation/100)^ActiveFairShares where ActiveFairShares is the number of active share identifiers.

For example, a computeReservation value of 50 indicates that Amazon Batch should reserve 50% of the maximum available VCPU if there is only one share identifier, 25% if there are two share identifiers, and 12.5% if there are three share identifiers. A computeReservation value of 25 indicates that Amazon Batch should reserve 25% of the maximum available VCPU if there is only one share identifier, 6.25% if there are two share identifiers, and 1.56% if there are three share identifiers.

- 7. In the **Share attributes** section, you can specify the share identifier and weight for each share identifier to associate with the fair-share scheduling policy.
  - a. Choose Add share identifier.
  - b. For **Share identifier**, specify the share identifier. If the string ends with '\*', this becomes a share identifier prefix used to match share identifiers for jobs. All of the share identifiers and share identifier prefixes in a scheduling policy must be unique and cannot overlap. For example, you can't have share identifiers prefix 'UserA\*' and share identifier 'UserA1' in the same fair-share scheduling policy.
  - c. For **Weight factor**, specify the relative weight for the share identifier. The default value is 1.0. A lower value has a higher priority for compute resources. If a share identifier prefix is used, jobs with share identifiers that start with the prefix will share the weight factor. This effectively increases the weight factor for those jobs, lowering their individual priority but maintaining the same weight factor for the share identifier prefix.
- 8. (Optional) In the **Tags** section, you can specify the key and value for each tag to associate with the scheduling policy. For more information, see <u>Tag your Amazon Batch resources</u>.
- 9. Choose **Submit** to finish and create your scheduling policy.

# Reference: Scheduling policy template

An empty scheduling policy template is shown below. You can use this template to create your scheduling policy which can then be saved to a file and used with the Amazon CLI --cli-inputjson option. For more information about these parameters, see CreateSchedulingPolicy in the Amazon Batch API Reference.



### Note

You can generate a job queue template with the following Amazon CLI command.

\$ aws batch create-scheduling-policy --generate-cli-skeleton

```
{
    "name": "",
    "fairsharePolicy": {
        "shareDecaySeconds": 0,
        "computeReservation": 0,
        "shareDistribution": [
             {
                 "shareIdentifier": "",
                 "weightFactor": 0.0
            }
        ]
    },
    "tags": {
        "KeyName": ""
    }
}
```

# Resource-aware scheduling

Amazon Batch schedules jobs based on the vCPU, GPU and the memory availability in the Compute Environment (CE) associated with the Job Queue (JQ). But sometimes, just the availability of these CE resources doesn't guarantee that the job will run successfully since it might be dependent on other required resources, so these jobs are cancelled or terminated. This results in inefficient use of compute resources. To solve this issue, resource-aware scheduling can check the availability of dependent, non-CE resources before it schedules the job to run on a CE.

Amazon Batch resource-aware scheduling lets you schedule jobs based on consumable resources that are needed to run your jobs— 3rd party license tokens, database access bandwidth, the need to throttle calls to a third-party API, and so on. You specify the consumable resources which are needed for a job to run, and Batch takes these resource dependencies into account when it schedules a job. You can avoid making manual interventions to eliminate job failures and long waits caused by a shortage of consumable resources. You can reduce the under-utilization of compute resources by allocating only the jobs that have all the required resources available.

Resource-aware scheduling is available for both FIFO and Fair-share scheduling policies and can be used with all compute platforms supported by Batch including EKS, ECS, and Fargate. It can be used with Array jobs, Multi-node parallel (MNP) jobs, and with regular Batch jobs.

To configure resource-aware scheduling, you first specify all the consumable resources needed to run your jobs, along with the total count available of each resource. Then, for each job that requires a consumable resource, you specify the name and required quantities of each resource needed. Batch keeps track of how many consumable resources are available for the jobs in your job queues and makes sure that a job is scheduled to run only when all the required consumable resources are available for the job to run successfully.

#### **Topics**

- Create consumable resources
- Specify the resources needed to run a job
- Check how many resources are in-use and available
- Update the quantity of a resource while it is in use by jobs
- Find the jobs that require a specific consumable resource
- Delete a consumable resource

### Create consumable resources

You must first create the consumable resources that represent the non-CE resources that are consumed when a job is running and are only available in limited quantities. Each consumable resource has a:

- resource name (consumableResourceName) that must be unique at the account level.
- (optional) resource type (resourceType) that indicates whether the resource is available to be re-used after a job completes. This can be one of:

Create consumable resources 152

- REPLENISHABLE (default)
- NON REPLENISHABLE
- total quantity (totalQuantity) that specifies the total amount of the consumable resource available.

The maximum number of consumable resources per account is 50k.

#### Console:

- 1. In the left navigation panel of the Amazon Batch console, choose Consumable resources.
- 2. Choose Create consumable resource.
- Enter a unique Resource name, the Total resource quantity, and select whether the Type of resource is Replenishable or Non-replenishable.
- 4. Choose Create consumable resource.

#### API:

Use the CreateConsumableResource API to define the resources you want.

# Specify the resources needed to run a job

When you register a job you can specify the name of one or more resources you created (consumableResource) and the quantity of that resource each instance of the job requires (quantity).

Batch keeps track of the available units of each resource at any given moment. For each job in the job queue, the Batch scheduler ensures that your job runs only when the specified resource dependencies are available.

If a consumable resource for the job is not available when the job reaches the head of the queue, the job will wait in RUNNABLE state until all the required resources become available or the job state time limit is reached (see <u>View job queue status</u>). Once Batch has validated that all the resources are available, the job transitions to the STARTING state and then to RUNNING. Resources are locked once the job moves to STARTING and are then unlocked when the job moves to SUCCEEDED or FAILED.

You can also update the quantity of a resource needed for a specific job when you submit the job.

Specify resources for a job 153

#### Console:

### To specify resources and their needed quantities when you define a job:

 Define a job using the job definition wizard from the <u>Amazon Batch console</u> (Job definitions -> Create).

- 2. In the wizard's Step 4: **Configure containers**, under **Consumable resource**, select the **Name** of a required resource from the list. In the **Requested value** field, enter the quantity of this resource needed by an instance of this job, then choose **Add consumable resource**.
- 3. Repeat the previous step for all the consumable resources required by the job. You can specify up to 5 resources for each job you define.
- 4. You'll see a list of the consumable resources you have created after you complete the job definition wizard but before you choose **Create job definition**.

#### To update needed quantities of resources when you submit a job:

- 1. In the left navigation pane of the <u>Amazon Batch console</u>, choose **Jobs**, then choose **Submit new job**.
- 2. In the wizard's Step 2: **Configure overrides**, under **Consumable resource overrides**, enter a new **Requested value** for any consumable resource whose needed quantity you want to override for the job.
- 3. After you have completed all the overrides you want to make for this job, choose **Next** to continue to **Review and submit**.

#### API:

When you register a job with the <u>RegisterJobDefinition API</u>, use the consumableResourceList in the consumableResourceProperties portion of the request to specify the consumable resources required to run an instance of the job, and the quantity of each.

When you submit a job with the <u>SubmitJob API</u> you can override the list of consumable resources and the quantity of each using the consumableResourcePropertiesOverride portion of the request. Note that this only overrides the quantity of the resource needed by each instance of the job, not the total quantity available.

Specify resources for a job 154

## Check how many resources are in-use and available

Batch lets you query the number of available resources (availableQuantity), the number of resources in use (inUseQuantity), and the total resources (totalQuantity) at a given moment.

Once a job goes to the STARTING state, the consumed resources will be subtracted from the available quantity of that resource. If the resource is REPLENISHABLE, the number of consumed resources will be added back to the available quantity as soon as the job has moved to either the SUCCEEDED or FAILED state, and the total quantity will remain the same. If the resource is NON\_REPLENISHABLE, the number of consumed resources is subtracted from both the total and available quantities and won't be added back whether the job moves to the SUCCEEDED or FAILED state.



#### Note

This information may lag by up to 30 seconds.

#### Console:

- In the left navigation panel of the Amazon Batch console, choose Consumable resources. 1.
- 2. Select either the **Replenishable** or **Non-replenishable** tab to view the resources of that type that you have created.
- For each **Replenishable** resource, the console displays the **Name**, the **Total** quantity of the resource, the number currently **In-use** and how many are still **Available**, along with a calculation of the Utilization (the number of resources in-use divided by the total quantity of that resource).

For each **Non-replenishable** resource, the console displays the **Name**, the number currently **In**use and how many are still Available.

You can also view current information about consumable resources from a job detail page in the console.

- In the left navigation panel of the Amazon Batch console, choose Jobs, then select a job's name to open the details page for that job.
- Information on both Replenishable resources and Non-replenishable resources are available to view if the job requires them. For both types, the console displays the resource's **Name**, the

155 Check resource usage

**Requested** quantity for the job, how many are still **Available**, the number currently **In-use**, the **Total** quantity of the resource, along with a calculation of the **Current utilization** (the number of resources in-use by the job divided by the total quantity of that resource).

#### API:

Use the DescribeConsumableResource API which returns the following information:

```
{
   "availableQuantity": number,
   "consumableResourceArn": "string",
   "createdAt": number,
   "inUseQuantity": number,
   "resourceType": "string",
   "tags": {
        "string" : "string"
   },
   "totalQuantity": number
}
```

The <u>ListConsumableResources API</u> also reports the number of resources in use (inUseQuantity) and the total number of resources currently available (totalQuantity) as part of its listing of all the consumable resources you have created in your account. This API also allows you filter the consumable resource list query based on the consumable resource name.

# Update the quantity of a resource while it is in use by jobs

You can reset the total quantity of a resource to a new value, add to the total quantity or subtract from it.

If the new total quantity you specify is greater than before, Batch schedules more jobs accordingly. If the new total quantity is less than before and there are no units of this resource in use, Batch just reduces the total (or available) quantity. If there are units in use, Batch reduces the available quantity immediately and, as jobs finish, Batch reduces the total (available) quantity so that it eventually arrives at the new number.

#### Console:

1. In the left navigation panel of the <u>Amazon Batch console</u>, choose **Consumable resources**.

2. Select either the **Replenishable** or **Non-replenishable** tab to view the resources of that type that you have created.

### 3. For **Replenishable** resources:

- Choose the resource you want to update, then select Actions and choose Set resources,
   Add resources, or Remove resources.
- 2. A pop-up window appears in which you can **Set total value**, **Add resources**, or **Remove resources** depending on which action you chose in the previous step. Enter the quantity you want to set as the new total value, you want to add to the total quantity or that you want to subtract from the total quantity, then select **Ok**.

### For **Non-replenishable** resources:

- 1. Choose the resource you want to update, then select **Actions** and choose **Set resources**, **Add resources**, or **Remove resources**.
- 2. A pop-up window appears in which you can **Set available value**, **Add resources**, or **Remove resources** depending on which action you chose in the previous step. Enter the quantity you want to set as the new available value, you want to add to the available quantity or that you want to subtract from the available quantity, then select **Ok**.

#### API:

Use the <u>UpdateConsumableResource API</u> to set a new total quantity for the resource, or to increase or reduce the total quantity.

# Find the jobs that require a specific consumable resource

Batch lets you retrieve a list of jobs that require a specific consumable resource.

#### Console:

- 1. In the left navigation panel of the Amazon Batch console, choose **Consumable resources**.
- 2. In the list, select the name of the consumable resource. The details page for that resource opens.
- 3. Under Search jobs, enter any filters you want to apply to the list of jobs. You can filter by the Job name ('equals', 'starts with'), the Date range (when the job was created), and by Additional criteria ('job queue', 'job definition', 'shared job identifier'). For each type of filter

you want to apply, select from the available options in the drop-down list and enter any additional information requested.

Choose **Search**.

4. A (filtered) list of jobs that require the consumable resource is displayed, including the job's name, status, number of requested units of the consumable resource, other needed consumable resources, and so on. Using this list, you can select one or more jobs to **Cancel** or **Terminate**. You can also select a job's name to open that job's detail page.

5. Under **Search jobs** you can now **Refresh results** or **Clear search** and start over.

#### **API:**

You can get a list of jobs that use a specific consumable resource with the <a href="ListJobsByConsumableResource API"><u>ListJobsByConsumableResource API</u></a>. This API also lets you filter the job list query using the job status or the job name.

### Delete a consumable resource

You can delete a consumable resource at any time, even when the jobs that require the resource are still running. Once a consumable resource is deleted, there may be a gap between the time the delete command is received and the job scheduler honors the delete, so it's possible jobs that consume the resource may be scheduled right after the delete call. If the deleted consumable resource has resource type (resourceType) REPLENISHABLE, this will be ignored when the jobs complete. If you delete a consumable resource and re-create it with the same name, it is considered to be the same resource and it can be used by RUNNABLE jobs.

#### Console:

- 1. In the left navigation panel of the <u>Amazon Batch console</u>, choose **Consumable resources**.
- 2. Select either the **Replenishable** or **Non-replenishable** tab to view the resources of that type that you have created.
- 3. Select each resource you want to delete, then choose **Delete**. A pop-up window **Delete consumable resource** appears. To confirm the deletion, choose **Delete**.

You can also delete a consumable resource from its detail page in the console.

1. In the left navigation panel of the Amazon Batch console, choose Consumable resources.

Delete a consumable resource 158

2. Select either the **Replenishable** or **Non-replenishable** tab to view the resources of that type that you have created.

3. Choose the name of the resource you want to delete. The details page of the consumable resource appears. Choose **Delete**. A pop-up window **Delete consumable resource** appears. To confirm the deletion, choose **Delete**.

### **API:**

Use the DeleteConsumableResource API to delete a consumable resource.

Delete a consumable resource 159

# Job definitions

Amazon Batch job definitions specify how jobs are to be run. While each job must reference a job definition, many of the parameters that are specified in the job definition can be overridden at runtime.

Some of the attributes specified in a job definition include:

- Which Docker image to use with the container in your job.
- How many vCPUs and how much memory to use with the container.
- The command the container should run when it is started.
- What (if any) environment variables should be passed to the container when it starts.
- Any data volumes that should be used with the container.
- What (if any) IAM role your job should use for Amazon permissions.

#### **Contents**

- Create a single-node job definition
- Create a multi-node parallel job definition
- Reference: Job definition template that uses ContainerProperties
- Create job definitions using EcsProperties
- Use the awslogs log driver
- Specify sensitive data
- Private registry authentication for jobs
- Amazon EFS volumes
- Reference: Job definition examples

# Create a single-node job definition

Before you can run jobs in Amazon Batch, you must create a job definition. This process varies slightly between single-node and multi-node parallel jobs. This topic covers specifically how to create a job definition for an Amazon Batch job that's not a multi-node parallel job (also known as gang scheduling).

You can create a multi-node parallel job definition on Amazon Elastic Container Service resources. For more information, see the section called "Create a multi-node parallel job definition".

### **Topics**

- Tutorial: Create a single-node job definition on Amazon EC2 resources
- Tutorial: Create a single-node job definition on Fargate resources
- Tutorial: Create a single-node job definition on Amazon EKS resources
- Tutorial: Create a single-node job definition with multiple containers on Amazon EC2 resources

# Tutorial: Create a single-node job definition on Amazon EC2 resources

Complete the following steps to create a single-node job definition on Amazon Elastic Compute Cloud (Amazon EC2) resources.

### To create a new job definition on Amazon EC2 resources:

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, choose the Amazon Web Services Region to use.
- 3. In the left navigation pane, choose **Job definitions**.
- 4. Choose Create.
- 5. For Orchestration type, choose Amazon Elastic Compute Cloud (Amazon EC2).
- 6. For **EC2 platform configuration**, turn off **Enable multi-node parallel** processing.
- 7. For **Name**, enter a unique name for your job definition. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 8. (Optional) For **Execution timeout**, enter the timeout value (in seconds). The execution timeout is the length of time before an unfinished job is terminated. If an attempt exceeds the timeout duration, the attempt is stopped and moves to a FAILED status. For more information, see <u>Job</u> timeouts. The minimum value is 60 seconds.
- 9. (Optional) Turn on **Scheduling priority**. Enter a scheduling priority value between 0 and 100. Higher values are given higher priority.
- 10. (Optional) For **Job attempts**, enter the number of times that Amazon Batch attempts to move the job to RUNNABLE status. Enter a number between 1 and 10.

11. (Optional) For Retry strategy conditions, choose Add evaluate on exit. Enter at least one parameter value and then choose an **Action**. For each set of conditions, **Action** must be set to either **Retry** or **Exit**. These actions mean the following:

- **Retry** Amazon Batch retries until the number of job attempts that you specified is reached.
- Exit Amazon Batch stops retrying the job.

#### Important

If you choose Add evaluate on exit, you must configure at least one parameter and either choose an Action or choose Remove evaluate on exit.

- 12. (Optional) Expand Tags and then choose Add tag to add tags to the resource. Enter a key and optional value, then choose **Add tag**.
- 13. (Optional) Turn on **Propagate tags** to propagate tags from the job and job definition to the Amazon ECS task.
- 14. Choose **Next page**.
- 15. In the **Container configuration** section:
  - For **Image**, choose the Docker image to use for your job. By default, images in the Docker a. Hub registry are available. You can also specify other repositories with repositoryurl/image: tag. The name can be up to 225 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), underscores (\_), colons (:), forward slashes (/), and number signs (#). This parameter maps to Image in the Create a container section of the Docker Remote API and the IMAGE parameter of docker run.



#### Note

Docker image architecture must match the processor architecture of the compute resources that they're scheduled on. For example, ARM-based Docker images can only run on ARM-based compute resources.

 Images in Amazon ECR Public repositories use the full registry/repository[:tag] or registry/repository[@digest] naming conventions (for example, public.ecr.aws/registry\_alias/my-web-app:latest).

 Images in Amazon ECR repositories use the full registry/ repository[:tag] naming convention (for example, aws\_account\_id.dkr.ecr.region.amazonaws.com/my-web-app:latest).

- Images in official repositories on Docker Hub use a single name (for example, ubuntu or mongo).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, amazon/amazon-ecs-agent).
- Images in other online repositories are qualified further by a domain name (for example, quay.io/assemblyline/ubuntu).
- For **Command**, enter the commands into the field as their **JSON** string array equivalent.

This parameter maps to Cmd in the Create a container section of the Docker Remote API and the COMMAND parameter to docker run. For more information about the Docker CMD parameter, see https://docs.docker.com/engine/reference/builder/#cmd.



#### Note

You can use default values for parameter substitution and placeholders in your command. For more information, see Parameters.

- (Optional) For **Execution role**, specify an IAM role that grants the Amazon ECS container c. agents permission to make Amazon API calls on your behalf. This feature uses Amazon ECS IAM roles for tasks. For more information, see Amazon ECS task execution IAM roles in the Amazon Elastic Container Service Developer Guide.
- For **Job Role configuration**, choose an IAM role that has permissions to the Amazon APIs. This feature uses Amazon ECS IAM roles for tasks. For more information, see IAM Roles for Tasks in the Amazon Elastic Container Service Developer Guide.



#### Note

Only roles that have the **Amazon Elastic Container Service Task Role** trust relationship are shown here. For more information about creating an IAM role for your Amazon Batch jobs, see Creating an IAM Role and Policy for your Tasks in the Amazon Elastic Container Service Developer Guide.

16. For **Parameters**, choose **Add parameters** to add parameter substitution placeholders as **Key** and optional Value pairs.

#### 17. In the **Environment configuration** section:

- For **vCPUs**, enter the number of vCPUs to reserve for the container. This parameter maps to CpuShares in the Create a container section of the Docker Remote API and the -cpu-shares option to docker run. Each vCPU is equivalent to 1,024 CPU shares. You must specify at least one vCPU.
- For **Memory**, enter the memory limit available to the container. If your container attempts to exceed the amount of memory that you specify here, the container is stopped. This parameter maps to Memory in the Create a container section of the Docker Remote API and the --memory option to docker run. You must specify at least 4 MiB of memory for a job.



### Note

To maximize your resource utilization, prioritize memory for jobs of a specific instance type. For more information, see Compute resource memory management.

- For **Number of GPUs**, choose the number of GPUs to reserve for the container. c.
- d. (Optional) For Environment variables, choose Add environment variable to add environment variables as name-value pairs. These variables are passed to the container.
- (Optional) For **Secrets**, choose **Add secret** to add secrets as a name-value e. pairs. These secrets are exposed in the container. For more information, see LogConfiguration:secretOptions.
- 18. Choose **Next page**.
- 19. In the **Linux configuration** section:
  - For **User**, enter the user name to use inside the container. This parameter maps to User in the Create a container section of the Docker Remote API and the --user option to docker run.
  - (Optional) To give the job container elevated permissions on the host instance (similar to the root user), drag the **Privileged** slider to the right. This parameter maps to Privileged in the Create a container section of the Docker Remote API and the -privileged option to docker run.

c. (Optional) Turn on **Enable init process** to run an init process inside the container. This process forwards signals and reaps processes.

### 20. (Optional) In the **Filesystem configuration** section:

- a. Turn on Enable read only filesystem to remove write access to the volume.
- b. For **Shared memory size**, enter the size (in MiB) of the /dev/shm volume.
- c. For **Max swap size**, enter the total amount of swap memory (in MiB) that the container can use.
- d. For **Swappiness** enter a value between 0 and 100 to indicate the swappiness behavior of the container. If you don't specify a value and swapping is enabled, the value defaults to 60. For more information, see LinuxParameters:swappiness.
- e. (Optional) Expand Additional configuration.
- f. (Optional) For **Tmpfs**, choose **Add tmpfs** to add a tmpfs mount.
- g. (Optional) For **Devices**, choose **Add device** to add a device:
  - i. For **Container path**, specify the path of in the container instance to expose the device mapped to the host instance. If you keep this blank, the host path is used in the container.
  - ii. For **Host path**, specify the path of a device in the host instance.
  - iii. For **Permissions**, choose one or more permissions to apply to the device. The available permissions are **READ**, **WRITE**, and **MKNOD**.
- h. (Optional) For **Volumes configuration**, choose **Add volume** to create a list of volumes to pass to the container. Enter **Name** and **Source path** for the volume and then choose **Add volume**. You can also choose to turn on **Enable EFS**.
- i. (Optional) For Mount points, choose Add mount points configuration to add mount points for data volumes. You must specify the source volume and container path. These mount points are passed to the Docker daemon on a container instance. You can also choose to make the volume Read only.
- j. (Optional) For **Ulimits configuration**, choose **Add ulimit** to add a ulimits value for the container. Enter **Name**, **Soft limit**, and **Hard limit** values, and then choose **Add ulimit**.

#### 21. In the **Task properties** section:

a. For **Execution role - conditional**, choose a role to allow Amazon ECS agents to make Amazon API calls on your behalf. For more information on creating an **Execution role**, see Tutorial: Create the IAM execution role.

Choose Enable ECS execute command, to enable access to the Amazon ECS container shell directly and bypass the host OS. You must choose a **Task role**.

#### Important

The **ECS** execute command requires that file system be writable.

- For **Task role**, choose an Amazon ECS Identity and Access Management (IAM) role to allow the container to make Amazon API calls on your behalf. For more information see, Amazon ECS task IAM role in the Amazon Elastic Container Service Developer Guide.
- 22. (Optional) In the **Logging configuration** section:
  - For **Log driver**, choose the log driver to use. For more information about the available log drivers, see LogConfiguration:logDriver.

### Note

By default, the awslogs log driver is used.

- For **Options**, choose **Add option** to add an option. Enter a name-value pair, and then choose **Add option**.
- For **Secrets**, choose **Add secret**. Enter a name-value pair and then choose **Add secret** to add a secret.

### (i) Tip

For more information, see LogConfiguration:secretOptions.

- 23. Choose Next page.
- 24. For **Job definition review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create job definition**.

## Tutorial: Create a single-node job definition on Fargate resources

Complete the following steps to create a single-node job definition on Amazon Fargate resources.

#### To create a new job definition on Fargate resources:

- 1. Open the Amazon Batch console at <a href="https://console.amazonaws.cn/batch/">https://console.amazonaws.cn/batch/</a>.
- 2. From the top navigation bar, choose the Amazon Web Services Region to use.
- 3. In the left navigation pane, choose **Job definitions**.
- 4. Choose Create.
- 5. For **Orchestration type**, choose **Fargate**. For more information, see <u>Fargate compute</u> environments.
- 6. For **Name**, enter a unique name for your job definition. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 7. (Optional) For **Execution timeout**, enter the timeout value (in seconds). The execution timeout is the length of time before an unfinished job is terminated. If an attempt exceeds the timeout duration, the attempt is stopped and moves to a FAILED status. For more information, see <u>Job timeouts</u>. The minimum value is 60 seconds.
- 8. (Optional) Turn on **Scheduling priority**. Enter a scheduling priority value between 0 and 100. Higher values are given higher priority over lower values.
- 9. (Optional) Expand **Tags**, and then choose **Add tag** to add tags to the resource. Turn on **Propagate tags** to propagate tags from the job and job definition.
- 10. In the Fargate platform configuration section:
  - a. For **Runtime platform**, choose the compute environment architecture.
  - For Operating System Family, choose the operating system for the compute environment.
  - c. For **CPU Architecture**, choose the vCPU architecture.
  - d. For **Fargate platform version**, enter LATEST or a specific runtime environment version.
  - e. (Optional) Turn on **Assign public IP** to assign a public IP address to a Fargate job network interface. For a job that's running in a private subnet to send outbound traffic to the internet, the private subnet requires a NAT gateway be attached to route requests to the internet. You might want to do this so that you can pull container images. For more information, see <a href="Amazon ECS"><u>Amazon ECS task networking</u></a> in the *Amazon Elastic Container Service Developer Guide*.

f. (Optional) For **Ephemeral storage**, enter the amount of ephemeral storage to allocate to the task. The amount of ephemeral storage must be between 21 GiB and 200 GiB. By default, 20 GiB of ephemeral storage is allocated if you don't enter a value.



#### (i) Note

Ephemeral storage requires Fargate platform version 1.4 or later.

- For **Execution role**, specify an IAM role that grants the Amazon ECS container and g. Fargate agents permission to make Amazon API calls on your behalf. This feature uses Amazon ECS IAM roles for task functionality. For more information including configuration prerequisites, see Amazon ECS task execution IAM roles in the Amazon Elastic Container Service Developer Guide.
- For **Job attempts**, enter the number of times that Amazon Batch attempts to move the job to a RUNNABLE status. Enter a number between 1 and 10.
- i. Optional) For **Retry strategy conditions**, choose **Add evaluate on exit**. Enter at least one parameter value and then choose an **Action**. For each set of conditions, **Action** must be set to either **Retry** or **Exit**. These actions mean the following:
  - Retry Amazon Batch retries until the number of job attempts that you specified is reached.
  - Exit Amazon Batch stops retrying the job.

#### Important

If you choose **Add evaluate on exit**, you must configure at least one parameter and choose an Action or choose Remove evaluate on exit.

- 11. Choose Next page.
- 12. In the **Container configuration** section:
  - For **Image**, choose the Docker image to use for your job. By default, images in the Docker Hub registry are available. You can also specify other repositories with repositoryurl/image: tag. The name can be up to 225 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), underscores (\_), colons (:), periods

(.), forward slashes (/), and number signs (#). This parameter maps to Image in the Create a container section of the Docker Remote API and the IMAGE parameter of docker run.



# Note

Docker image architecture must match the processor architecture of the compute resources that they're scheduled on. For example, ARM-based Docker images can only run on ARM-based compute resources.

- Images in Amazon ECR Public repositories use the full registry/repository[:tag] or registry/repository[@digest] naming conventions (for example, public.ecr.aws/registry\_alias/my-web-app:latest).
- Images in Amazon ECR repositories use the full registry/ repository[:tag] naming convention (for example, aws\_account\_id.dkr.ecr.region.amazonaws.com/my-web-app:latest).
- Images in official repositories on Docker Hub use a single name (for example, ubuntu or mongo).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, amazon/amazon-ecs-agent).
- Images in other online repositories are qualified further by a domain name (for example, quay.io/assemblyline/ubuntu).
- For **Command**, enter the commands into the field as their JSON string array equivalent. b.

This parameter maps to Cmd in the Create a container section of the Docker Remote API and the COMMAND parameter to docker run. For more information about the Docker CMD parameter, see https://docs.docker.com/engine/reference/builder/#cmd.



### Note

You can use default values for parameter substitution and placeholders in your command. For more information, see Parameters.

(Optional) Add parameters to the job definition as name-value mappings to override the c. job definition defaults. To add a parameter:

For Parameters, choose Add parameters, enter a name-value pair, then choose Add parameter.

# Important

If you choose **Add parameter**, you must either configure at least one parameter or choose Remove parameter

- In the **Environment configuration** section:
  - i. For **Job role configuration**, choose an IAM role that has permissions to the Amazon APIs. This feature uses Amazon ECS IAM roles for task functionality. For more information, see IAM Roles for Tasks in the Amazon Elastic Container Service Developer Guide.

## Note

Only roles that have the **Amazon Elastic Container Service Task Role** trust relationship are shown here. For more information about how to create an IAM role for your Amazon Batch jobs, see Creating an IAM Role and Policy for your Tasks in the Amazon Elastic Container Service Developer Guide.

- For **vCPUs**, enter the number of vCPUs to reserve for the container. This parameter ii. maps to CpuShares in the Create a container section of the Docker Remote API and the --cpu-shares option to docker run. Each vCPU is equivalent to 1,024 CPU shares. You must specify at least one vCPU.
- iii. For **Memory**, enter the memory limit that's available to the container. If your container attempts to exceed the memory specified here, the container is stopped. This parameter maps to Memory in the Create a container section of the Docker Remote API and the --memory option to docker run. You must specify at least 4 MiB of memory for a job.

If you use GuardDuty Runtime Monitoring, there is a slight memory overhead for the GuardDuty security agent. Therefore the memory limit must include the size of the GuardDuty security agent. For information about the GuardDuty security agent memory limits, see CPU and memory limits in the GuardDuty User Guide. For

> information about the best practices, see How do I remediate out of memory errors on my Fargate tasks after enabling Runtime Monitoring in the Amazon ECS Developer Guide.

### (i) Note

To maximize your resource utilization, prioritize memory for jobs of a specific instance type. For more information, see Compute resource memory management.

- (Optional) For Environment variables, choose Add environment variable to add environment variables as name-value pairs. These variables are passed to the container.
- f. (Optional) For **Secrets**, choose **Add secret** to add secrets as a name-value pairs. These secrets are exposed in the container. For more information, see LogConfiguration:secretOptions.
- g. Choose **Next page**.
- 13. (Optional) In the **Linux configuration** section:
  - For **User**, enter a user name to use inside the container. a.
  - Turn on **Enable init process** to run an init process inside the container. This process forwards signals and reaps processes.
  - Turn on **Enable read only filesystem** to remove write access to the volume. C.
  - d. (Optional) Expand Additional configuration.
  - For Mount points configuration, choose Add mount points configuration to add mount e. points for data volumes. You must specify the source volume and container path. These mount points are passed to the Docker daemon on a container instance.
  - For **Volumes configuration**, choose **Add volume** to create a list of volumes to pass to the f. container. Enter a **Name** and **Source path** for the volume, and then choose **Add volume**.
  - In the **Task properties** section:
    - For **Execution role conditional**, choose a role to allow Amazon ECS agents to make Amazon API calls on your behalf. For more information on creating an **Execution role**, see Tutorial: Create the IAM execution role.
    - Choose **Enable ECS execute command**, to enable access to the Amazon ECS container ii. shell directly and bypass the host OS. You must choose a **Task role**.

## Important

The **ECS** execute command requires that file system be writable.

iii. For Task role, choose an Amazon ECS Identity and Access Management (IAM) role to allow the container to make Amazon API calls on your behalf. For more information see, Amazon ECS task IAM role in the Amazon Elastic Container Service Developer Guide.

# In the **Logging configuration** section:

i. (Optional) For **Log driver**, choose the log driver to use. For more information about the available log drivers, see LogConfiguration:logDriver.



# Note

By default, the awslogs log driver is used.

- (Optional) For **Options**, choose **Add option** to add an option. Enter a name-value pair, ii. and then choose Add option.
- (Optional) For **Secrets**, choose **Add secret** to add a secret. Then, enter a name-value pair, and choose Add secret.



# 🚺 Tip

For more information, see LogConfiguration:secretOptions.

- 14. Choose Next page.
- 15. For **Job definition review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create job definition**.

# Tutorial: Create a single-node job definition on Amazon EKS resources

Complete the following steps to create a single-node job definition on Amazon Elastic Kubernetes Service (Amazon EKS).

## To create a new job definition on Amazon EKS resources:

- Open the Amazon Batch console at https://console.amazonaws.cn/batch/. 1.
- 2. From the top navigation bar, choose the Amazon Web Services Region to use.
- 3. In the left navigation pane, choose **Job definitions**.
- Choose Create. 4.
- For Orchestration type, choose Elastic Kubernetes Service (EKS). 5.
- 6. For **Name**, enter a unique name for your job definition. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores ().
- (Optional) For **Execution timeout**, enter the timeout value (in seconds). The execution timeout is the length of time before an unfinished job is terminated. If an attempt exceeds the timeout duration, the attempt is stopped and moves to a FAILED status. For more information, see Job timeouts. The minimum value is 60 seconds.
- (Optional) Turn on **Scheduling priority**. Enter a scheduling priority value between 0 and 100. Higher values are given higher priority over lower values.
- 9. (Optional) Expand **Tags**, and then choose **Add tag** to add tags to the resource.
- 10. Choose **Next page**.
- 11. In the **EKS pod properties** section:
  - For **Service account name**, enter an account that provides an identity for processes that a. run in a pod.
  - Turn **Host network** on to use the Kubernetes pod network model and open a listening port for incoming connections. Turn this setting off for outgoing communications only.
  - For **DNS policy**, choose one of the following:
    - No value (null) The pod ignores the DNS settings from the Kubernetes environment.
    - **Default** The pod inherits the name resolution configuration from the node that it runs on.



### Note

If a DNS policy isn't specified, **Default** isn't the default DNS policy. Instead, ClusterFirst is used.

• ClusterFirst – Any DNS guery that doesn't match the configured cluster domain suffix is forwarded to the upstream nameserver that's inherited from the node.

- ClusterFirstWithHostNet Use if Host network is turned on.
- (Optional) For **Volumes**, select **Add volume**, then:
  - i. Add a **Name** for your volume.
  - ii. (Optional) Add the **Host path** for the directory on the host.
  - iii. (Optional) Add a Medium and a Size limit to configure a Kubernetes emptyDir.
  - (Optional) Provide a **Secret name** for the pod and whether the secret is **Optional**. iv.
  - (Optional) Define a Claim name to attach a Kubernetes Persistent Volume Claim to ٧. the pod, and whether it is **Read only**.
- (Optional) For **Pod labels**, choose **Add pod labels**, then enter a name-value pair. e.

# Important

The prefix for a pod label can't contain kubernetes.io/, k8s.io/, or batch.amazonaws.com/.

f. (Optional) For **Pod annotations**, choose **Add annotations**, then enter a name-value pair.

## 

The prefix for a pod annotation can't contain kubernetes.io/, k8s.io/, or batch.amazonaws.com/.

- Choose Next page. g.
- In the **Container configuration** section:
  - i. For Name, enter a unique name for the container. The name must start with a letter or number, and can be up to 63 characters long. It can contain uppercase and lowercase letters, numbers, and hyphens (-).
  - For **Image**, choose the Docker image to use for your job. By default, images in the Docker Hub registry are available. You can also specify other repositories with repository-url/image:tag. The name can be up to 255 characters long. It can contain uppercase and lowercase letters, numbers, hyphens (-), underscores (\_), colons (:), periods (.), forward slashes (/), and number signs (#). This parameter maps to

> Image in the Create a container section of the Docker Remote API and the IMAGE parameter of docker run

# Note

Docker image architecture must match the processor architecture of the compute resources that they're scheduled on. For example, ARM-based Docker images can only run on ARM-based compute resources.

- Images in Amazon ECR Public repositories use the full registry/ repository[:tag] or registry/repository[@digest] naming conventions (for example, public.ecr.aws/registry\_alias/my-web-app:latest).
- Images in Amazon ECR repositories use the full registry/ repository[:tag] naming convention (for example, aws\_account\_id.dkr.ecr.region.amazonaws.com/my-web-app:latest).
- Images in official repositories on Docker Hub use a single name (for example, ubuntu or mongo).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, amazon/amazon-ecs-agent).
- Images in other online repositories are qualified further by a domain name (for example, quay.io/assemblyline/ubuntu).
- iii. (Optional) For **Image pull policy**, choose when images are pulled.
- iv. (Optional) For **Command**, enter a JSON command to pass to the container.
- (Optional) For **Arguments**, enter arguments to pass to the container. If an argument isn't provided, the container image command is used.
- (Optional) You can add parameters to the job definition as name-value mappings to i. override the job definition defaults. To add a parameter:
  - For **Parameters**, enter a name-value pair, then choose **Add parameter**.



### Important

If you choose Add parameter, you must configure at least one parameter or choose Remove parameter

#### In the **Environment configuration** section: j.

For vCPUs, enter the number of vCPUs to reserve for the container. This parameter i. maps to CpuShares in the Create a container section of the Docker Remote API and the --cpu-shares option to docker run. Each vCPU is equivalent to 1,024 CPU shares. You must specify at least one vCPU.

For **Memory**, enter the memory limit available to the container. If your container attempts to exceed the memory specified here, the container is stopped. This parameter maps to Memory in the Create a container section of the Docker Remote API and the --memory option to docker run. You must specify at least 4 MiB of memory for a job.



### Note

To maximize your resource utilization, prioritize memory for jobs of a specific instance type. For more information, see Compute resource memory management.

- (Optional) For Environment variables, choose Add environment variable to add k. environment variables as name-value pairs. These variables are passed to the container.
- l. (Optional) For Volume mount:
  - Choose Add volume mount. i.
  - Enter a Name, and then enter a Mount path in the container where the volume is mounted. Enter a **SubPath** to specify a sub-path inside the referenced volume instead of its root.
  - iii. Choose **Read only** to remove write permissions to the volume.
  - Choose Add volume mount.
- m. (Optional) For **Run as user**, enter a user ID to run the container process.



## Note

The user ID must exist in the image for the container to run.

(Optional) For **Run as group**, enter a group ID to run the container process runtime.



# Note

The group ID must exist in the image for the container to run.

(Optional) To give your job's container elevated permissions on the host instance (similar to the root user), drag the **Privileged** slider to the right. This parameter maps to Privileged in the Create a container section of the Docker Remote API and the -privileged option to docker run.

- p. (Optional) Turn on **Read-only root filesystem** to remove write access to the root filesystem.
- (Optional) Turn on **Run as non-root** to run the containers in the pod as a non-root user.



# Note

If **Run as non-root** is turned on, the kubelet validates the image at runtime to verify that image doesn't run as UID 0.

- r. Choose **Next page**.
- 12. For **Job definition review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create job definition**.

# Tutorial: Create a single-node job definition with multiple containers on Amazon EC2 resources

Complete the following steps to create a single-node job definition with multiple containers on Amazon Elastic Compute Cloud (Amazon EC2) resources.

# To create a new job definition on Amazon EC2 resources:

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, choose the Amazon Web Services Region to use.
- 3. In the left navigation pane, choose **Job definitions**.
- 4. Choose **Create**.
- 5. For Orchestration type, choose Amazon Elastic Compute Cloud (Amazon EC2).
- For Job definition structure, turn off Use legacy containerProperties structure processing. 6.

- For EC2 platform configuration, turn off Enable multi-node parallel processing. 7.
- 8. Choose **Next**.
- In **General configuration** section, enter the following: 9.
  - For Name, enter a unique name for your job definition. The name can be up to 128 a. characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores ( ).
  - For **Execution timeout optional**, enter the timeout value (in seconds). The execution timeout is the length of time before an unfinished job is terminated. If an attempt exceeds the timeout duration, the attempt is stopped and moves to a FAILED status. For more information, see Job timeouts. The minimum value is 60 seconds.
  - Turn on **Scheduling priority optional**. Enter a scheduling priority value between 0 and 100. Higher values are given higher priority.
  - Expand Tags optional and then choose Add tag to add tags to the resource. Enter a key and optional value, then choose Add tag.
  - Turn on **Propagate tags** to propagate tags from the job and job definition to the Amazon ECS task.
- 10. In **Retry strategy optional** section, enter the following:
  - a. For **Job attempts**, enter the number of times that Amazon Batch attempts to move the job to RUNNABLE status. Enter a number between 1 and 10.
  - b. For **Retry strategy conditions**, choose **Add evaluate on exit**. Enter at least one parameter value and then choose an Action. For each set of conditions, Action must be set to either **Retry** or **Exit**. These actions mean the following:
    - Retry Amazon Batch retries until the number of job attempts that you specified is reached.
    - Exit Amazon Batch stops retrying the job.

## Important

If you choose **Add evaluate on exit**, you must configure at least one parameter and either choose an Action or choose Remove evaluate on exit.

11. In **Task properties** section, enter the following:

For **Execution role** - conditional, choose a role to allow Amazon ECS agents to make a. Amazon API calls on your behalf. For more information on creating an **Execution role**, see Tutorial: Create the IAM execution role.

b. Choose **Enable ECS execute command**, to enable access to the Amazon ECS container shell directly and bypass the host OS. You must choose a **Task role**.

## 

The **ECS** execute command requires that file system be writable.

- For **Task role**, choose an Amazon ECS Identity and Access Management (IAM) role to allow c. the container to make Amazon API calls on your behalf. For more information see, Amazon ECS task IAM role in the Amazon Elastic Container Service Developer Guide.
- For IPC mode choose host, task, or none. If host is specified, then all the containers that are within the tasks that specified the host IPC mode on the same container instance share the same IPC resources with the host Amazon EC2 instance. If task is specified, all the containers that are within the specified task share the same IPC resources. If none is specified, then IPC resources within the containers of a task are private and not shared with other containers in a task or on the container instance. If no value is specified, then the IPC resource namespace sharing depends on the Docker daemon setting on the container instance.
- For **PID mode** choose host or task. For example, monitoring sidecars might need pidMode to access information about other containers running in the same task. If host is specified, all containers within the tasks that specified the host PID mode on the same container instance share the same process namespace with the host Amazon EC2 instance. If task is specified, all containers within the specified task share the same process namespace. If no value is specified, the default is a private namespace for each container.
- 12. In the **Consumable resource** section, enter the following:
  - Enter a unique Name and the Requested value. a.
  - You can add more consumable resources by choosing **Add consumable resource**.
- 13. In the **Storage** section, enter the following:
  - Enter a Name and Source path for the volume and then choose Add volume. You can also a. choose to turn on Enable EFS.

- You can add more Volumes by choosing **Add volume**.
- 14. For Parameters, choose Add parameters to add parameter substitution placeholders as Key and optional **Value** pairs.
- 15. Choose **Next page**.
- 16. In the **Container configuration** section:
  - For **Name**, enter a name for the container. a.
  - For **Essential container**, enable if the container is essential. b.
  - For Image, choose the Docker image to use for your job. By default, images in the Docker c. Hub registry are available. You can also specify other repositories with repositoryurl/image: tag. The name can be up to 225 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), underscores (\_), colons (:), forward slashes (/), and number signs (#). This parameter maps to Image in the Create a container section of the Docker Remote API and the IMAGE parameter of docker run.

# Note

Docker image architecture must match the processor architecture of the compute resources that they're scheduled on. For example, ARM-based Docker images can only run on ARM-based compute resources.

- Images in Amazon ECR Public repositories use the full registry/repository[:tag] or registry/repository[@digest] naming conventions (for example, public.ecr.aws/registry\_alias/my-web-app:latest).
- Images in Amazon ECR repositories use the full registry/ repository[:tag] naming convention (for example, aws\_account\_id.dkr.ecr.region.amazonaws.com/my-web-app:latest).
- Images in official repositories on Docker Hub use a single name (for example, ubuntu or mongo).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, amazon/amazon-ecs-agent).
- Images in other online repositories are qualified further by a domain name (for example, quay.io/assemblyline/ubuntu).
- For **Resource requirements** configure each of the following:

- For **vCPUs**, choose the number of CPUs for the container. i.
- ii. For **Memory**, choose the amount of memory for the container.
- For **GPU** *optional*, choose the number of GPUs for the container.
- For **User**, enter the user name to use inside the container. e.
- f. Turn on **Enable read only filesystem** to remove write access to the volume.
- Turn on **Privileged** to give the job container elevated permissions on the host instance, g. similar to the root user.
- For **Command**, enter the commands into the field as their **JSON** string array equivalent.

This parameter maps to Cmd in the Create a container section of the Docker Remote API and the COMMAND parameter to docker run. For more information about the Docker CMD parameter, see https://docs.docker.com/engine/reference/builder/#cmd.



# Note

You can use default values for parameter substitution and placeholders in your command. For more information, see Parameters.

- For **Repository credentials optional** enter the ARN of the secret containing your i. credentials.
- j. For **Environment variables - optional**, choose **Add environment variables** to add environment variables to pass to the container.
- k. In the **Linux parameters - optional** section:
  - i. Turn on **Enable init process** to run an init process inside the container.
  - ii. For **Shared memory size**, enter the size (in MiB) of the /dev/shm volume
  - iii. For Max swap size, enter the total amount of swap memory (in MiB) that the container can use.
  - iv. For **Swappiness** enter a value between 0 and 100 to indicate the swappiness behavior of the container. If you don't specify a value and swapping is enabled, the value defaults to 60.
  - For **Devices**, choose **Add device** to add a device:

> A. For **Container path**, specify the path of in the container instance to expose the device mapped to the host instance. If you keep this blank, the host path is used in the container.

- B. For **Host path**, specify the path of a device in the host instance.
- C. For **Permissions**, choose one or more permissions to apply to the device. The available permissions are **READ**, **WRITE**, and **MKNOD**.

vi. For **Tmpfs**, choose **Add tmpfs** to add a tmpfs mount.

l.

# Note

Firelens logging has to be done in a dedicated container. To configure Firelens logging:

- In every container, except your dedicated firelens container, set the Logging driver to awsfirelens
- In your Firelens container set the **Firelens Configuration optional** and the **Logging configuration - optional** to the logging destination

# In the Firelens Configuration - optional section:

### Important

Amazon Batch enforces host network mode on non-MNP, non-FARGATE Amazon ECS jobs. Root user is required for Amazon ECS Firelens. When running tasks that use the host network mode, Amazon ECS advises against running containers using the root user (UID 0) for better security. Therefore, all non-MNP, non-FARGATE ECS jobs with Firelens logging will not meet security best practice.

- i. For **Type**, choose either fluentd or fluentbit.
- For **Options**, enter the name/value pair of the option. You can add more **Options** ii. using Added option.
- m. In the **Logging configuration** *optional* section:

For **Log driver**, choose the log driver to use. For more information about the available i. log drivers, see LogConfiguration:logDriver.



# Note

By default, the awslogs log driver is used.

- For **Options**, choose **Add option** to add an option. Enter a name-value pair, and then ii. choose **Add option**.
- iii. For Secrets, choose Add secret. Enter a name-value pair and then choose Add secret to add a secret.



# (i) Tip

For more information, see LogConfiguration:secretOptions.

- For Mount points optional, choose Add mount points to add mount points for data volumes. You must specify the source volume and container path.
- For **Secrets optional**, choose **Add secret** to add a secret. Then, enter a name-value pair, and choose Add secret.



For more information, see LogConfiguration:secretOptions.

- For **Ulimits optional**, choose **Add ulimit** to add a ulimits value for the container. Enter Name, Soft limit, and Hard limit values, and then choose Add ulimit.
- For **Dependencies optional**, choose **Add container dependencies**. Choose the name of the container and it's state to determine when this container starts.
- 17. If you only have one container configured then you must choose Add container and complete configuring the new container. Otherwise, choose **Next** to review.

# Create a multi-node parallel job definition

Before you can run jobs in Amazon Batch, you must create a job definition. This process varies slightly between single-node and multi-node parallel jobs. This topic covers specifically how

to create a job definition for an Amazon Batch multi-node parallel job (also known as gang scheduling). For more information, see Multi-node parallel jobs.



# Note

Amazon Fargate doesn't support multi-node parallel jobs.

#### **Contents**

• Tutorial: Create a multi-node parallel job definition on Amazon EC2 resources

# Tutorial: Create a multi-node parallel job definition on Amazon EC2 resources

To create a multi-node parallel job definition on Amazon Elastic Compute Cloud (Amazon EC2) resources.



# Note

To create a single-node job definition, see Tutorial: Create a single-node job definition on Amazon EC2 resources.

# To create a multi-node parallel job definition on Amazon EC2 resources:

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Job definitions**.
- 4. Choose Create.
- 5. For Orchestration type, choose Amazon Elastic Compute Cloud (Amazon EC2).
- 6. For **Enable multi-node parallel**, turn on multi-node parallel.
- For **Name**, enter a unique name for your job definition. The name can be up to 128 characters 7. long, and can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).

(Optional) For **Execution timeout**, specify the maximum number of seconds that you want job attempts to run. If an attempt exceeds the timeout duration, the attempt is stopped and moves to a FAILED status. For more information, see Job timeouts.

- 9. (Optional) Turn on **Scheduling priority**. Enter a scheduling priority value between 0 and 100. Higher values are given higher priority over lower values.
- 10. (Optional) For **Job attempts**, enter the number of times that Amazon Batch attempts to move the job to RUNNABLE status. Enter a number between 1 and 10.
- 11. (Optional) For **Retry strategy conditions**, choose **Add evaluate on exit**. Enter at least one parameter value and then choose an Action. For each set of conditions, Action must be set to either **Retry** or **Exit**. These actions mean the following:
  - **Retry** Amazon Batch retries until the number of job attempts that you specified is reached.
  - Exit Amazon Batch stops retrying the job.

# Important

If you choose Add evaluate on exit, you must configure at least one parameter and either choose an Action or choose Remove evaluate on exit.

- 12. (Optional) Expand Tags and then choose Add tag to add tags to the resource. Enter a key and optional value, and then choose **Add tag**. You can also turn on **Propagate tags** to propagate tags from the job and job definition to the Amazon ECS task.
- 13. Choose **Next page**.
- 14. For **Number of nodes**, enter the total number of nodes to use for your job.
- 15. For Main node, enter the node index to use for the main node. The default main node index is 0.
- 16. For **Instance type**, choose an instance type.



# Note

The instance type that you choose applies to all nodes.

- 17. For Parameters, choose Add parameters to add parameter substitution placeholders as Key and optional Value pairs.
- 18. In the **Node ranges** section:

- Select **Add node range**. This creates a **Node range** section. a.
- b. For **Target nodes**, specify the range for your node group, using range\_start:range\_end notation.

You can create up to five node ranges for the nodes that you specified for your job. Node ranges use the index value for a node, and the node index begins at 0. Make sure that range end index value of your final node group is one less than the number of nodes that you specified. For example, suppose that you specified 10 nodes, and you want to use a single node group. Then, your end range is 9.

For **Image**, choose the Docker image to use for your job. By default, images in the Docker Hub registry are available. You can also specify other repositories with repositoryurl/image: tag. The name can be up to 225 characters long. It can contain uppercase and lowercase letters, numbers, hyphens (-), underscores (\_), colons (:), forward slashes (/), and number signs (#). This parameter maps to Image in the Create a container section of the Docker Remote API and the IMAGE parameter of **docker run**.

# Note

Docker image architecture must match the processor architecture of the compute resources that they're scheduled on. For example, ARM-based Docker images can only run on ARM-based compute resources.

- Images in Amazon ECR Public repositories use the full registry/repository[:tag] or registry/repository[@digest] naming conventions (for example, public.ecr.aws/registry\_alias/my-web-app:latest).
- Images in Amazon ECR repositories use the full registry/ repository[:tag] naming convention. For example, aws\_account\_id.dkr.ecr.region.amazonaws.com/my-web-app:latest
- Images in official repositories on Docker Hub use a single name (for example, ubuntu or mongo).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, amazon/amazon-ecs-agent).
- Images in other online repositories are qualified further by a domain name (for example, quay.io/assemblyline/ubuntu).

For **Command**, enter the commands into the field as their **JSON** string array equivalent. d.

This parameter maps to Cmd in the Create a container section of the Docker Remote API and the COMMAND parameter to docker run. For more information about the Docker CMD parameter, see https://docs.docker.com/engine/reference/builder/#cmd.



## Note

You can use default values for parameter substitution and placeholders in your command. For more information, see Parameters.

- For **vCPUs**, specify the number of vCPUs to reserve for the container. This parameter maps e. to CpuShares in the Create a container section of the Docker Remote API and the -cpu-shares option to docker run. Each vCPU is equivalent to 1,024 CPU shares. You must specify at least one vCPU.
- f. For **Memory**, specify the hard limit (in MiB) of memory to present to the job's container. If your container attempts to exceed the memory specified here, the container is stopped. This parameter maps to Memory in the Create a container section of the Docker Remote API and the --memory option to docker run. You must specify at least 4 MiB of memory for a job.



## Note

To maximize your resource utilization, you can provide your jobs as much memory as possible for a particular instance type. For more information, see Compute resource memory management.

- (Optional) For **Number of GPUs**, specify the number of GPUs your job uses. The job runs g. on a container with the specified number of GPUs that are pinned to that container.
- (Optional) For Job role, you can specify an IAM role that provides the container in your job with permissions to use the Amazon APIs. This feature uses Amazon ECS IAM roles for task functionality. For more information including configuration prerequisites, see IAM Roles for Tasks in the Amazon Elastic Container Service Developer Guide.



# Note

For jobs that are running on Fargate resources, a job role is required.



## Note

Only roles that have the Amazon Elastic Container Service Task Role trust relationship are shown here. For more information about creating an IAM role for your Amazon Batch jobs, see Creating an IAM Role and Policy for your Tasks in the Amazon Elastic Container Service Developer Guide.

(Optional) For Execution role, specify an IAM role that grants the Amazon ECS container i. agents permission to make Amazon API calls on your behalf. This feature uses Amazon ECS IAM roles for task functionality. For more information, see Amazon ECS task execution IAM roles in the Amazon Elastic Container Service Developer Guide.

## 19. (Optional) Expand Additional configuration:

- For Environment variables, choose Add environment variable to add environment variables as name-value pairs. These variables are passed to the container.
- For **Job role configuration**, you can specify an IAM role that provides the container in your job with permissions to use the Amazon APIs. This feature uses Amazon ECS IAM roles for task functionality. For more information including configuration prerequisites, see IAM Roles for Tasks in the Amazon Elastic Container Service Developer Guide.



#### Note

For jobs that are running on Fargate resources, a job role is required.



# Note

Only roles that have the Amazon Elastic Container Service Task Role trust relationship are shown here. For more information about how to create an IAM role for your Amazon Batch jobs, see Creating an IAM Role and Policy for your Tasks in the Amazon Elastic Container Service Developer Guide.

For **Execution role**, specify an IAM role that grants the Amazon ECS container agents c. permission to make Amazon API calls on your behalf. This feature uses Amazon ECS IAM roles for task functionality. For more information, see Amazon ECS task execution IAM roles in the Amazon Elastic Container Service Developer Guide.

## 20. In the **Security Configuration** section:

a. (Optional) To give your job's container elevated privileges on the host instance (similar to the root user), turn on **Privileged**. This parameter maps to Privileged in the <u>Create a</u> container section of the <u>Docker Remote API</u> and the --privileged option to **docker run**.

- b. (Optional) For **User**, enter the user name to use inside the container. This parameter maps to User in the <u>Create a container</u> section of the <u>Docker Remote API</u> and the --user option to **docker run**.
- c. (Optional) For Secrets, choose Add secret to add secrets as a name-value pairs. These secrets are exposed in the container. For more information, see LogConfiguration:secretOptions.

# 21. In the **Linux configuration** section:

- a. Turn on **Enable read only filesystem** to remove write access to the volume.
- b. (Optional) Turn on **Enable init process** to run an init process inside the container. This process forwards signals and reaps processes.
- c. For **Shared memory size**, enter the size (in MiB) of the /dev/shm volume.
- d. For **Max swap size**, enter the total amount of swap memory (in MiB) that the container can use.
- e. For **Swappiness** enter a value between 0 and 100 to indicate the swappiness behavior of the container. If you don't specify a value and swapping is enabled, value defaults to 60. For more information, see LinuxParameters:swappiness.
- f. (Optional) For **Devices**, choose **Add device** to add a device:
  - i. For **Container path**, specify the path of in the container instance to expose the device mapped to the host instance. If you keep this blank, the host path is used in the container.
  - ii. For **Host path**, specify the path of a device in the host instance.
  - iii. For **Permissions**, choose one or more permissions to apply to the device. The available permissions are **READ**, **WRITE**, and **MKNOD**.
- 22. (Optional) For **Mount points**, choose **Add mount points configuration** to add mount points for data volumes. You must specify the source volume and container path. These mount points are passed to the Docker daemon on a container instance. You can also choose to make the volume **Read only**.

23. (Optional) For **Ulimits configuration**, choose **Add ulimit** to add a ulimits value for the container. Enter Name, Soft limit, and Hard limit values, and then choose Add ulimit.

- 24. (Optional) For Volumes configuration, choose Add volume to create a list of volumes to pass to the container. Enter Name and Source path for the volume and then choose Add volume. You can also choose to turn on **Enable EFS**.
- 25. (Optional) For **Tmpfs**, choose **Add tmpfs** to add a tmpfs mount.
- 26. In the **Task properties** section:
  - For **Execution role conditional**, choose a role to allow Amazon ECS agents to make Amazon API calls on your behalf. For more information on creating an **Execution role**, see Tutorial: Create the IAM execution role.

b.

### Important

To use **ECS execute command** your compute environment must meet the compute environment considerations for multi node parallel jobs.

Choose **Enable ECS execute command**, to enable access to the Amazon ECS container shell directly and bypass the host OS. You must choose a **Task role**.



## 

The **ECS** execute command requires that file system be writable.

- For **Task role**, choose an Amazon ECS Identity and Access Management (IAM) role to allow the container to make Amazon API calls on your behalf. For more information see, Amazon ECS task IAM role in the Amazon Elastic Container Service Developer Guide.
- 27. (Optional) In the **Logging configuration** section:
  - For **Log driver**, choose the log driver to use. For more information about the available log a. drivers, see LogConfiguration:logDriver.



## Note

By default, the awslogs log driver is used.

For **Options**, choose **Add option** to add an option. Enter a name-value pair, and then choose Add option.

For **Secrets**, choose **Add secret**. Enter a name-value pair and then choose **Add secret** to add a secret.



# (i) Tip

For more information, see LogConfiguration:secretOptions.

- 28. Choose Next page.
- 29. For **Job definition review**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create job definition**.

# Reference: Job definition template that uses **ContainerProperties**

The following is an empty job definition template that includes a single container. You can use this template to create your job definition, which can then be saved to a file and used with the Amazon CLI --cli-input-json option. For more information about these parameters, see JobDefinition.

# Note

You can generate a single-container job definition template with the following Amazon CLI command:

```
$ aws batch register-job-definition --generate-cli-skeleton
```

```
{
    "jobDefinitionName": "",
    "type": "container",
    "parameters": {
        "KeyName": ""
    },
    "schedulingPriority": 0,
    "containerProperties": {
        "image": "",
```

```
"vcpus": 0,
"memory": 0,
"command": [
    11.11
],
"jobRoleArn": "",
"executionRoleArn": "",
"volumes": [
    {
        "host": {
            "sourcePath": ""
        },
        "name": "",
        "efsVolumeConfiguration": {
            "fileSystemId": "",
            "rootDirectory": "",
            "transitEncryption": "ENABLED",
            "transitEncryptionPort": 0,
            "authorizationConfig": {
                "accessPointId": "",
                "iam": "DISABLED"
            }
        }
    }
],
"environment": [
    {
        "name": "",
        "value": ""
    }
],
"mountPoints": [
    {
        "containerPath": "",
        "readOnly": true,
        "sourceVolume": ""
    }
],
"readonlyRootFilesystem": true,
"privileged": true,
"ulimits": [
    {
        "hardLimit": 0,
        "name": "",
```

```
"softLimit": 0
    }
],
"user": "",
"instanceType": "",
"resourceRequirements": [
    {
        "value": "",
        "type": "MEMORY"
    }
],
"linuxParameters": {
    "devices": [
        {
            "hostPath": "",
            "containerPath": "",
            "permissions": [
                "WRITE"
            ]
        }
    ],
    "initProcessEnabled": true,
    "sharedMemorySize": 0,
    "tmpfs": [
        {
            "containerPath": "",
            "size": 0,
            "mountOptions": [
                11 11
            ]
        }
    ],
    "maxSwap": 0,
    "swappiness": 0
},
"logConfiguration": {
    "logDriver": "syslog",
    "options": {
        "KeyName": ""
   },
    "secretOptions": [
        {
            "name": "",
            "valueFrom": ""
```

```
}
        ]
    },
    "secrets": [
        {
            "name": "",
            "valueFrom": ""
    ],
    "networkConfiguration": {
        "assignPublicIp": "DISABLED"
    },
    "fargatePlatformConfiguration": {
        "platformVersion": ""
    }
},
"nodeProperties": {
    "numNodes": 0,
    "mainNode": 0,
    "nodeRangeProperties": [
        {
            "targetNodes": "",
            "container": {
                "image": "",
                "vcpus": 0,
                "memory": 0,
                "command": [
                     11 11
                ],
                "jobRoleArn": "",
                "executionRoleArn": "",
                "volumes": [
                     {
                         "host": {
                             "sourcePath": ""
                         },
                         "name": "",
                         "efsVolumeConfiguration": {
                             "fileSystemId": "",
                             "rootDirectory": "",
                             "transitEncryption": "DISABLED",
                             "transitEncryptionPort": 0,
                             "authorizationConfig": {
                                 "accessPointId": "",
```

```
"iam": "ENABLED"
            }
        }
    }
],
"environment": [
    {
        "name": "",
        "value": ""
    }
],
"mountPoints": [
    {
        "containerPath": "",
        "readOnly": true,
        "sourceVolume": ""
    }
],
"readonlyRootFilesystem": true,
"privileged": true,
"ulimits": [
    {
        "hardLimit": 0,
        "name": "",
        "softLimit": 0
    }
],
"user": "",
"instanceType": "",
"resourceRequirements": [
    {
        "value": "",
        "type": "MEMORY"
    }
],
"linuxParameters": {
    "devices": [
        {
            "hostPath": "",
            "containerPath": "",
            "permissions": [
                "WRITE"
        }
```

```
],
                     "initProcessEnabled": true,
                     "sharedMemorySize": 0,
                     "tmpfs": [
                         {
                             "containerPath": "",
                             "size": 0,
                             "mountOptions": [
                         }
                     ],
                     "maxSwap": 0,
                     "swappiness": 0
                },
                "logConfiguration": {
                     "logDriver": "awslogs",
                     "options": {
                         "KeyName": ""
                     },
                     "secretOptions": [
                             "name": "",
                             "valueFrom": ""
                     ]
                },
                "secrets": [
                     {
                         "name": "",
                         "valueFrom": ""
                     }
                ],
                "networkConfiguration": {
                     "assignPublicIp": "DISABLED"
                },
                "fargatePlatformConfiguration": {
                     "platformVersion": ""
                }
            }
        }
    ]
},
"retryStrategy": {
```

```
"attempts": 0,
    "evaluateOnExit": [
        {
            "onStatusReason": "",
            "onReason": "",
            "onExitCode": "",
            "action": "RETRY"
        }
    ]
},
"propagateTags": true,
"timeout": {
    "attemptDurationSeconds": 0
},
"tags": {
    "KeyName": ""
"platformCapabilities": [
    "EC2"
],
"eksProperties": {
    "podProperties": {
        "serviceAccountName": "",
        "hostNetwork": true,
        "dnsPolicy": "",
        "containers": [
            {
                 "name": "",
                 "image": "",
                 "imagePullPolicy": "",
                 "command": [
                     11 11
                 ],
                 "args": [
                    11 11
                 ],
                 "env": [
                     {
                         "name": "",
                         "value": ""
                     }
                 ],
                 "resources": {
                     "limits": {
```

```
"KeyName": ""
                         },
                         "requests": {
                             "KeyName": ""
                         }
                     },
                     "volumeMounts": [
                         {
                             "name": "",
                             "mountPath": "",
                             "readOnly": true
                         }
                     ],
                     "securityContext": {
                         "runAsUser": 0,
                         "runAsGroup": 0,
                         "privileged": true,
                         "readOnlyRootFilesystem": true,
                         "runAsNonRoot": true
                     }
                }
            ],
            "volumes": [
                 {
                     "name": "",
                     "hostPath": {
                         "path": ""
                     },
                     "emptyDir": {
                         "medium": "",
                         "sizeLimit": ""
                     },
                     "secret": {
                         "secretName": "",
                         "optional": true
                     }
                }
            ]
        }
    }
}
```

# Reference: Job definition parameters for ContainerProperties

Job definitions that use ContainerProperties are split into several parts:

- the job definition name
- the type of the job definition
- the parameter substitution placeholder defaults
- the container properties for the job
- the Amazon EKS properties for the job definition that are necessary for jobs run on Amazon EKS resources
- the node properties that are necessary for a multi-node parallel job
- the platform capabilities that are necessary for jobs run on Fargate resources
- the default tag propagation details of the job definition
- the default retry strategy for the job definition
- the default scheduling priority for the job definition
- the default tags for the job definition
- · the default timeout for the job definition

#### **Contents**

- Job definition name
- Type
- Parameters
- Container properties
- Amazon EKS properties
- Platform capabilities
- Propagate tags
- Node properties
- Retry strategy
- Scheduling priority
- Tags

#### Timeout

## Job definition name

jobDefinitionName

When you register a job definition, you specify a name. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_). The first job definition that's registered with that name is given a revision of 1. Any subsequent job definitions that are registered with that name are given an incremental revision number.

Type: String

Required: Yes

# **Type**

type

When you register a job definition, you specify the type of job. If the job runs on Fargate resources, then multinode isn't supported. For more information about multi-node parallel jobs, see Create a multi-node parallel job definition.

Type: String

Valid values: container | multinode

Required: Yes

# **Parameters**

parameters

When you submit a job, you can specify parameters that replace the placeholders or override the default job definition parameters. Parameters in job submission requests take precedence over the defaults in a job definition. This means that you can use the same job definition for multiple jobs that use the same format. You can also programmatically change values in the command at submission time.

Type: String to string map

Required: No

When you register a job definition, you can use parameter substitution placeholders in the command field of a job's container properties. The syntax is as follows.

```
"command": [
    "ffmpeg",
    "-i",
    "Ref::inputfile",
    "-c",
    "Ref::codec",
    "-o",
    "Ref::outputfile"
]
```

In the above example, there are <code>Ref::inputfile</code>, <code>Ref::codec</code>, and <code>Ref::outputfile</code> parameter substitution placeholders in the command. You can use the parameters object in the job definition to set default values for these placeholders. For example, to set a default for the <code>Ref::codec</code> placeholder, you specify the following in the job definition:

```
"parameters" : {"codec" : "mp4"}
```

When this job definition is submitted to run, the *Ref::codec* argument in the command for the container is replaced with the default value, mp4.

# **Container properties**

When you register a job definition, specify a list of container properties that are passed to the Docker daemon on a container instance when the job is placed. The following container properties are allowed in a job definition. For single-node jobs, these container properties are set at the job definition level. For multi-node parallel jobs, container properties are set in the <u>Node properties</u> level, for each node group.

command

The command that's passed to the container. This parameter maps to Cmd in the <u>Create a</u> container section of the <u>Docker Remote API</u> and the COMMAND parameter to **docker run**. For

more information about the Docker CMD parameter, see https://docs.docker.com/engine/ reference/builder/#cmd.

```
"command": ["string", ...]
```

Type: String array

Required: No

environment

The environment variables to pass to a container. This parameter maps to Env in the Create a container section of the Docker Remote API and the --env option to docker run.

# Important

We don't recommend that you use plaintext environment variables for sensitive information, such as credential data.

## Note

Environment variables must not start with AWS\_BATCH. This naming convention is reserved for variables that are set by the Amazon Batch service.

Type: Array of key-value pairs

Required: No

name

The name of the environment variable.

Type: String

Required: Yes, when environment is used.

value

The value of the environment variable.

Type: String

Required: Yes, when environment is used.

```
"environment" : [
    { "name" : "envName1", "value" : "envValue1" },
    { "name" : "envName2", "value" : "envValue2" }
]
```

### executionRoleArn

When you register a job definition, you can specify an IAM role. The role provides the Amazon ECS container agent with permissions to call the API actions that are specified in its associated policies on your behalf. Jobs that run on Fargate resources must provide an execution role. For more information, see Amazon Batch IAM execution role.

Type: String

Required: No

fargatePlatformConfiguration

The platform configuration for jobs that run on Fargate resources. Jobs that run on EC2 resources must not specify this parameter.

Type: FargatePlatformConfiguration object

Required: No

platformVersion

The Amazon Fargate platform version use for the jobs, or LATEST to use a recent, approved version of the Amazon Fargate platform.

Type: String

Default: LATEST

Required: No

image

The image used to start a job. This string is passed directly to the Docker daemon. Images in the Docker Hub registry are available by default. You can also specify other repositories

with repository-url/image: tag. Up to 255 letters (uppercase and lowercase), numbers, hyphens, underscores, colons, periods, forward slashes, and number signs are allowed. This parameter maps to Image in the Create a container section of the Docker Remote API and the IMAGE parameter of docker run.



# Note

Docker image architecture must match the processor architecture of the compute resources that they're scheduled on. For example, ARM-based Docker images can only run on ARM-based compute resources.

- Images in Amazon ECR Public repositories use the full registry/repository[:tag] or registry/repository[@digest] naming conventions (for example, public.ecr.aws/registry\_alias/my-web-app:latest).
- Images in Amazon ECR repositories use the full registry/repository: [tag] naming convention. For example, aws\_account\_id.dkr.ecr.region.amazonaws.com/my-webapp:latest.
- Images in official repositories on Docker Hub use a single name (for example, ubuntu or mongo).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, amazon/amazon-ecs-agent).
- Images in other online repositories are qualified further by a domain name (for example, quay.io/assemblyline/ubuntu).

Type: String

Required: Yes

### instanceType

The instance type to use for a multi-node parallel job. All node groups in a multi-node parallel job must use the same instance type. This parameter isn't valid for single-node container jobs or for jobs that run on Fargate resources.

Type: String

Required: No

## jobRoleArn

When you register a job definition, you can specify an IAM role. The role provides the job container with permissions to call the API actions that are specified in its associated policies on your behalf. For more information, see <a href="IAM Roles for Tasks">IAM Roles for Tasks</a> in the Amazon Elastic Container Service Developer Guide.

Type: String

Required: No

### linuxParameters

Linux-specific modifications that are applied to the container, such as details for device mappings.

```
"linuxParameters": {
    "devices": [
        {
            "hostPath": "string",
            "containerPath": "string",
            "permissions": [
                 "READ", "WRITE", "MKNOD"
            ]
        }
    ],
    "initProcessEnabled": true|false,
    "sharedMemorySize": 0,
    "tmpfs": [
        {
            "containerPath": "string",
            "size": integer,
            "mountOptions": [
                 "string"
            ]
        }
    ],
    "maxSwap": integer,
    "swappiness": integer
}
```

Type: LinuxParameters object

Required: No

#### devices

List of devices mapped into the container. This parameter maps to Devices in the Create a container section of the Docker Remote API and the --device option to docker run.



### Note

This parameter isn't applicable to jobs that run on Fargate resources.

Type: Array of Device objects

Required: No

hostPath

Path where the device available in the host container instance is.

Type: String

Required: Yes

containerPath

Path where the device is exposed in the container is. If this isn't specified, the device is exposed at the same path as the host path.

Type: String

Required: No

permissions

Permissions for the device in the container. If this isn't specified the permissions are set to READ, WRITE, and MKNOD.

Type: Array of strings

Required: No

Valid values: READ | WRITE | MKNOD

## initProcessEnabled

If true, run an init process inside the container that forwards signals and reaps processes. This parameter maps to the --init option to docker run. This parameter requires version

1.25 of the Docker Remote API or greater on your container instance. To check the Docker Remote API version on your container instance, log into your container instance and run the following command: sudo docker version | grep "Server API version"

Type: Boolean

Required: No

maxSwap

The total amount of swap memory (in MiB) a job can use. This parameter is translated to the --memory-swap option to docker run where the value is the sum of the container memory plus the maxSwap value. For more information, see --memory-swap details in the Docker documentation.

If a maxSwap value of 0 is specified, the container doesn't use swap. Accepted values are 0 or any positive integer. If the maxSwap parameter is omitted, the container uses the swap configuration for the container instance that it runs on. A maxSwap value must be set for the swappiness parameter to be used.



### Note

This parameter isn't applicable to jobs that run on Fargate resources.

Type: Integer

Required: No

sharedMemorySize

The value for the size (in MiB) of the /dev/shm volume. This parameter maps to the --shmsize option to docker run.



# Note

This parameter isn't applicable to jobs that run on Fargate resources.

Type: Integer

Required: No

#### swappiness

You can use this to tune a container's memory swappiness behavior. A swappiness value of 0 causes swapping to not happen unless absolutely necessary. A swappiness value of 100 causes pages to be swapped aggressively. Accepted values are whole numbers between 0 and 100. If the swappiness parameter isn't specified, a default value of 60 is used. If a value isn't specified for maxSwap, then this parameter is ignored. If maxSwap is set to 0, the container doesn't use swap. This parameter maps to the --memory-swappiness option to docker run.

Consider the following when you use a per-container swap configuration.

• Swap space must be enabled and allocated on the container instance for the containers to use.



# Note

The Amazon ECS optimized AMIs don't have swap enabled by default. You must enable swap on the instance to use this feature. For more information, see Instance Store Swap Volumes in the Amazon EC2 User Guide or How do I allocate memory to work as swap space in an Amazon EC2 instance by using a swap file?.

- The swap space parameters are only supported for job definitions using EC2 resources.
- If the maxSwap and swappiness parameters are omitted from a job definition, each container has a default swappiness value of 60. The total swap usage is limited to two times the memory reservation of the container.



#### Note

This parameter isn't applicable to jobs that run on Fargate resources.

Type: Integer

Required: No

tmpfs

The container path, mount options, and size of the tmpfs mount.

Type: Array of Tmpfs objects



# Note

This parameter isn't applicable to jobs that run on Fargate resources.

Required: No

containerPath

The absolute file path in the container where the tmpfs volume is mounted.

Type: String

Required: Yes

mountOptions

The list of tmpfs volume mount options.

```
Valid values: "defaults" | "ro" | "rw" | "suid" | "nosuid" | "dev" | "nodev" | "exec" |
"noexec"|"sync"|"async"|"dirsync"|"remount"|"mand"|"nomand"|"atime"
| "noatime" | "diratime" | "nodiratime" | "bind" | "rbind" | "unbindable" |
"runbindable" | "private" | "rprivate" | "shared" | "rshared" | "slave" | "rslave"
| "relatime" | "norelatime" | "strictatime" | "nostrictatime" | "mode" | "uid" |
"gid"|"nr_inodes"|"nr_blocks"|"mpol"
```

Type: Array of strings

Required: No

size

The size (in MiB) of the tmpfs volume.

Type: Integer

Required: Yes

logConfiguration

The log configuration specification for the job.

This parameter maps to LogConfig in the Create a container section of the Docker Remote API and the --log-driver option to docker run. By default, containers use the same logging driver that the Docker daemon uses. However, the container can use a different logging

driver than the Docker daemon by specifying a log driver with this parameter in the container definition. To use a different logging driver for a container, the log system must be either configured on the container instance or on another log server to provide remote logging options. For more information about the options for different supported log drivers, see Configure logging drivers in the Docker documentation.



#### Note

Amazon Batch currently supports a subset of the logging drivers available to the Docker daemon (shown in the LogConfiguration data type).

This parameter requires version 1.18 of the Docker Remote API or greater on your container instance. To check the Docker Remote API version on your container instance, log into your container instance and run the following command: sudo docker version | grep "Server API version"

```
"logConfiguration": {
    "devices": [
        {
            "logDriver": "string",
            "options": {
                 "optionName1" : "optionValue1",
                 "optionName2" : "optionValue2"
            }
            "secretOptions": [
              {
                   "name" : "secretOptionName1",
                   "valueFrom" : "secretOptionArn1"
              },
              {
                   "name" : "secretOptionName2",
                   "valueFrom" : "secretOptionArn2"
              }
            ]
        }
    ]
}
```

Type: LogConfiguration object

## Required: No

# logDriver

The log driver to use for the job. By default, Amazon Batch enables the awslogs log driver. The valid values that are listed for this parameter are log drivers that the Amazon ECS container agent can communicate with by default.

This parameter maps to LogConfig in the Create a container section of the Docker Remote API and the --log-driver option to docker run. By default, jobs use the same logging driver that the Docker daemon uses. However, the job can use a different logging driver than the Docker daemon by specifying a log driver with this parameter in the job definition. If you want to specify another logging driver for a job, the log system must be configured on the container instance in the compute environment. Or, alternatively, configure it on another log server to provide remote logging options. For more information about the options for different supported log drivers, see Configure logging drivers in the Docker documentation.



#### Note

Amazon Batch currently supports a subset of the logging drivers that are available to the Docker daemon. Additional log drivers might be available in future releases of the Amazon ECS container agent.

The supported log drivers are awslogs, fluentd, gelf, json-file, journald, logentries, syslog, and splunk.



### Note

Jobs that run on Fargate resources are restricted to the awslogs and splunk log drivers.

This parameter requires version 1.18 of the Docker Remote API or greater on your container instance. To check the Docker Remote API version on your container instance, log into your container instance and run the following command: sudo docker version | grep "Server API version"



# Note

The Amazon ECS container agent that runs on a container instance must register the logging drivers that are available on that instance with the ECS AVAILABLE LOGGING DRIVERS environment variable. Otherwise, the containers placed on that instance can't use these log configuration options. For more information, see Amazon ECS Container Agent Configuration in the Amazon Elastic Container Service Developer Guide.

## awslogs

Specifies the Amazon CloudWatch Logs logging driver. For more information, see Use the awslogs log driver and Amazon CloudWatch Logs logging driver in the Docker documentation.

#### fluentd

Specifies the Fluentd logging driver. For more information including usage and options, see Fluentd logging driver in the Docker documentation.

# gelf

Specifies the Graylog Extended Format (GELF) logging driver. For more information including usage and options, see Graylog Extended Format logging driver in the Docker documentation.

# journald

Specifies the journald logging driver. For more information including usage and options, see Journald logging driver in the Docker documentation.

# json-file

Specifies the JSON file logging driver. For more information including usage and options, see JSON File logging driver in the Docker documentation.

# splunk

Specifies the Splunk logging driver. For more information including usage and options, see Splunk logging driver in the Docker documentation.

# syslog

Specifies the syslog logging driver. For more information including usage and options, see Syslog logging driver in the Docker documentation.

Type: String

Required: Yes

Valid values: awslogs | fluentd | gelf | journald | json-file | splunk | syslog



### Note

If you have a custom driver that's not listed earlier that you would like to work with the Amazon ECS container agent, you can fork the Amazon ECS container agent project that's available on GitHub and customize it to work with that driver. We encourage you to submit pull requests for changes that you want to have included. However, Amazon Web Services doesn't currently support requests that run modified copies of this software.

# options

Log configuration options to send to a log driver for the job.

This parameter requires version 1.19 of the Docker Remote API or greater on your container instance.

Type: String to string map

Required: No

secretOptions

An object that represents the secret to pass to the log configuration. For more information, see Specify sensitive data.

Type: object array

Required: No

name

The name of the log driver option to set in the job.

Type: String

Required: Yes

valueFrom

The Amazon Resource Name (ARN) of the secret to expose to the log configuration of the container. The supported values are either the full ARN of the Secrets Manager secret or the full ARN of the parameter in the SSM Parameter Store.



### Note

If the SSM Parameter Store parameter exists in the same Amazon Web Services Region as the task that you're launching, then you can use either the full ARN or name of the parameter. If the parameter exists in a different Region, then the full ARN must be specified.

Type: String

Required: Yes

memory

This parameter is deprecated, use resourceRequirements instead.

The number of MiB of memory reserved for the job.

As an example for how to use resourceRequirements, if your job definition contains syntax that's similar to the following.

```
"containerProperties": {
  "memory": 512
}
```

The equivalent syntax using resourceRequirements is as follows.

```
"containerProperties": {
 "resourceRequirements": [
   {
      "type": "MEMORY",
      "value": "512"
```

```
}
]
}
```

Type: Integer

Required: Yes

mountPoints

The mount points for data volumes in your container. This parameter maps to Volumes in the Create a container section of the Docker Remote API and the --volume option to docker run.

Type: Object array

Required: No

sourceVolume

The name of the volume to mount.

Type: String

Required: Yes, when mountPoints is used.

containerPath

The path on the container where to mount the host volume.

Type: String

Required: Yes, when mountPoints is used.

readOnly

If this value is true, the container has read-only access to the volume. If this value is false, then the container can write to the volume.

Type: Boolean

Required: No

Default: False

networkConfiguration

The network configuration for jobs that run on Fargate resources. Jobs that run on EC2 resources must not specify this parameter.

```
"networkConfiguration": {
    "assignPublicIp": "string"
}
```

Type: Object array

Required: No

assignPublicIp

Indicates whether the job has a public IP address. This is required if the job needs outbound network access.

Type: String

Valid values: ENABLED | DISABLED

Required: No

Default: DISABLED

privileged

When this parameter is true, the container is given elevated permissions on the host container instance (similar to the root user). This parameter maps to Privileged in the <u>Create a container</u> section of the <u>Docker Remote API</u> and the --privileged option to <u>docker run</u>. This parameter isn't applicable to jobs that run on Fargate resources. Don't provide it or specify it as false.

```
"privileged": true|false
```

Type: Boolean

Required: No

# readonlyRootFilesystem

When this parameter is true, the container is given read-only access to its root file system. This parameter maps to ReadonlyRootfs in the <u>Create a container</u> section of the <u>Docker Remote</u> API and the --read-only option to **docker run**.

```
"readonlyRootFilesystem": true|false
```

Type: Boolean

Required: No

resourceRequirements

The type and amount of a resource to assign to a container. The supported resources include GPU, MEMORY, and VCPU.

```
"resourceRequirements" : [
    {
        "type": "GPU",
        "value": "number"
    }
]
```

Type: Object array

Required: No

type

The type of resource to assign to a container. The supported resources include GPU, MEMORY, and VCPU.

Type: String

Required: Yes, when resourceRequirements is used.

value

The quantity of the specified resource to reserve for the container. The values vary based on the type specified.

# type="GPU"

The number of physical GPUs to reserve for the container. The number of GPUs reserved for all containers in a job cannot exceed the number of available GPUs on the compute resource that the job is launched on.

# type="MEMORY"

The hard limit (in MiB) of memory to present to the container. If your container attempts to exceed the memory specified here, the container is killed. This parameter maps to Memory in the Create a container section of the Docker Remote API and the --memory option to docker run. You must specify at least 4 MiB of memory for a job. This is required but can be specified in several places for multi-node parallel (MNP) jobs. It must be specified for each node at least once. This parameter maps to Memory in the Create a container section of the Docker Remote API and the --memory option to docker run.



#### Note

If you're trying to maximize your resource utilization by providing your jobs as much memory as possible for a particular instance type, see Compute resource memory management.

For jobs that run on Fargate resources, then value must match one of the supported values. Moreover, the VCPU values must be one of the values that's supported for that memory value.

VCPU	MEMORY
0.25 vCPU	512, 1024, and 2048 MiB
0.5 vCPU	1024-4096 MiB in 1024 MiB increments
1 vCPU	2048-8192 MiB in 1024 MiB increments
2 vCPU	4096-16384 MiB in 1024 MiB increments
4 vCPU	8192-30720 MiB in 1024 MiB increments
8 vCPU	16384-61440 MiB in 4096 MiB increments

VCPU	MEMORY
16 vCPU	32768-122880 MiB in 8192 MiB increments

type="VCPU"

The number of vCPUs reserved for the job. This parameter maps to CpuShares in the Create a container section of the Docker Remote API and the --cpu-shares option to docker run. Each vCPU is equivalent to 1,024 CPU shares. For jobs that run on EC2 resources, you must specify at least one vCPU. This is required but can be specified in several places. It must be specified for each node at least once.

For jobs that run on Fargate resources, value must match one of the supported values and the MEMORY values must be one of the values that's supported for that VCPU value. The supported values are 0.25, 0.5, 1, 2, 4, 8, and 16.

The default for the Fargate On-Demand vCPU resource count quota is 6 vCPUs. For more information about Fargate quotas, see <u>Amazon Fargate quotas</u> in the *Amazon Web Services General Reference*.

Type: String

Required: Yes, when resourceRequirements is used.

#### secrets

The secrets for the job that are exposed as environment variables. For more information, see Specify sensitive data.

Type: Object array

Required: No

name

The name of the environment variable that contains the secret.

Type: String

Required: Yes, when secrets is used.

#### valueFrom

The secret to expose to the container. The supported values are either the full Amazon Resource Name (ARN) of the Secrets Manager secret or the full ARN of the parameter in the SSM Parameter Store.



### (i) Note

If the SSM Parameter Store parameter exists in the same Amazon Web Services Region as the job you're launching, then you can use either the full ARN or name of the parameter. If the parameter exists in a different Region, then the full ARN must be specified.

Type: String

Required: Yes, when secrets is used.

#### ulimits

A list of ulimits values to set in the container. This parameter maps to Ulimits in the Create a container section of the Docker Remote API and the --ulimit option to docker run.

```
"ulimits": [
    "name": string,
    "softLimit": integer,
    "hardLimit": integer
  }
]
```

Type: Object array

Required: No

name

The type of the ulimit.

Type: String

Required: Yes, when ulimits is used.

hardLimit

The hard limit for the ulimit type.

Type: Integer

Required: Yes, when ulimits is used.

softLimit

The soft limit for the ulimit type.

Type: Integer

Required: Yes, when ulimits is used.

user

The user name to use inside the container. This parameter maps to User in the <u>Create a</u> container section of the <u>Docker Remote API</u> and the --user option to **docker run**.

```
"user": "string"
```

Type: String

Required: No

vcpus

This parameter is deprecated, use resourceRequirements instead.

The number of vCPUs reserved for the container.

As an example for how to use resourceRequirements, if your job definition contains lines similar to this:

```
"containerProperties": {
   "vcpus": 2
}
```

The equivalent lines using resourceRequirements is as follows.

```
"containerProperties": {
    "resourceRequirements": [
        {
            "type": "VCPU",
            "value": "2"
        }
     ]
}
```

Type: Integer

Required: Yes

volumes

When you register a job definition, you can specify a list of volumes that are passed to the Docker daemon on a container instance. The following parameters are allowed in the container properties:

```
"transitEncryptionPort": number
    }
  }
]
```

#### name

The name of the volume. Up to 255 letters (uppercase and lowercase), numbers, hyphens, and underscores are allowed. This name is referenced in the sourceVolume parameter of container definition mountPoints.

Type: String

Required: No

#### host

The contents of the host parameter determine whether your data volume persists on the host container instance and where it's stored. If the host parameter is empty, then the Docker daemon assigns a host path for your data volume. However, the data isn't guaranteed to persist after the container associated with it stops running.



# Note

This parameter isn't applicable to jobs that run on Fargate resources.

Type: Object

Required: No

sourcePath

The path on the host container instance that's presented to the container. If this parameter is empty, then the Docker daemon assigns a host path for you.

If the host parameter contains a sourcePath file location, then the data volume persists at the specified location on the host container instance until you delete it manually. If the sourcePath value doesn't exist on the host container instance, the Docker daemon creates it. If the location does exist, the contents of the source path folder are exported.

Type: String

Required: No

efsVolumeConfiguration

This parameter is specified when you're using an Amazon Elastic File System file system for task storage. For more information, see Amazon EFS volumes.

Type: Object

Required: No

authorizationConfig

The authorization configuration details for the Amazon EFS file system.

Type: String

Required: No

accessPointId

The Amazon EFS access point ID to use. If an access point is specified, the root directory value that's specified in the EFSVolumeConfiguration must either be omitted or set to /. This enforces the path that's set on the EFS access point. If an access point is used, transit encryption must be enabled in the EFSVolumeConfiguration. For more information, see <a href="Working with Amazon EFS">Working with Amazon EFS</a> Access Points in the Amazon Elastic File System User Guide.

Type: String

Required: No

iam

Determines whether to use the Amazon Batch job IAM role defined in a job definition when mounting the Amazon EFS file system. If enabled, transit encryption must be enabled in the EFSVolumeConfiguration. If this parameter is omitted, the default value of DISABLED is used. For more information, see <u>Use Amazon EFS access points</u>.

Type: String

Valid values: ENABLED | DISABLED

Required: No

# fileSystemId

The Amazon EFS file system ID to use.

Type: String

Required: No

rootDirectory

The directory within the Amazon EFS file system to mount as the root directory inside the host. If this parameter is omitted, the root of the Amazon EFS volume is used. If you specify /, it has the same effect as omitting this parameter. The maximum length is 4,096 characters.

### 

If an EFS access point is specified in the authorizationConfig, the root directory parameter must either be omitted or set to /. This enforces the path that's set on the Amazon EFS access point.

Type: String

Required: No

transitEncryption

Determines whether to enable encryption for Amazon EFS data in transit between the Amazon ECS host and the Amazon EFS server. Transit encryption must be enabled if Amazon EFS IAM authorization is used. If this parameter is omitted, the default value of DISABLED is used. For more information, see Encrypting data in transit in the Amazon Elastic File System User Guide.

Type: String

Valid values: ENABLED | DISABLED

Required: No

transitEncryptionPort

The port to use when sending encrypted data between the Amazon ECS host and the Amazon EFS server. If you don't specify a transit encryption port, it uses the port

selection strategy that the Amazon EFS mount helper uses. The value must be between 0 and 65,535. For more information, see <u>EFS Mount Helper</u> in the *Amazon Elastic File System User Guide*.

Type: Integer

Required: No

# **Amazon EKS properties**

An object with various properties that are specific to Amazon EKS based jobs. This must not be specified for Amazon ECS based job definitions.

podProperties

The properties for the Kubernetes pod resources of a job.

Type: EksPodProperties object

Required: No

containers

The properties of the container that's used on the Amazon EKS pod.

Type: EksContainer object

Required: No

args

An array of arguments to the entrypoint. If this isn't specified, the CMD of the container image is used. This corresponds to the args member in the <a href="Entrypoint">Entrypoint</a> portion of the <a href="Pod">Pod</a> in Kubernetes. Environment variable references are expanded using the container's environment.

If the referenced environment variable doesn't exist, the reference in the command isn't changed. For example, if the reference is to "\$(NAME1)" and the NAME1 environment variable doesn't exist, the command string will remain "\$(NAME1)." \$\$ is replaced with \$, and the resulting string isn't expanded. For example, \$\$(VAR\_NAME) is passed as \$(VAR\_NAME) whether or not the VAR\_NAME environment variable exists. For more

information, see CMD in the Dockerfile reference and Define a command and arguments for a pod in the Kubernetes documentation.

Type: Array of strings

Required: No

command

The entrypoint for the container. This isn't run within a shell. If this isn't specified, the ENTRYPOINT of the container image is used. Environment variable references are expanded using the container's environment.

If the referenced environment variable doesn't exist, the reference in the command isn't changed. For example, if the reference is to "\$(NAME1)" and the NAME1 environment variable doesn't exist, the command string will remain "\$(NAME1)." \$\$ is replaced with \$ and the resulting string isn't expanded. For example, \$\$(VAR\_NAME) will be passed as \$(VAR\_NAME) whether or not the VAR\_NAME environment variable exists. The entrypoint can't be updated. For more information, see ENTRYPOINT in the Dockerfile reference and Define a command and arguments for a container and Entrypoint in the Kubernetes documentation.

Type: Array of strings

Required: No

env

The environment variables to pass to a container.



# Note

Environment variables cannot start with "AWS\_BATCH". This naming convention is reserved for variables that Amazon Batch sets.

Type: Array of EksContainerEnvironmentVariable objects

Required: No

name

The name of the environment variable.

Type: String

Required: Yes

value

The value of the environment variable.

Type: String

Required: No

image

The Docker image used to start the container.

Type: String

Required: Yes

imagePullPolicy

The image pull policy for the container. Supported values are Always, IfNotPresent, and Never. This parameter defaults to IfNotPresent. However, if the :latest tag is specified, it defaults to Always. For more information, see <a href="Updating images">Updating images</a> in the Kubernetes documentation.

Type: String

Required: No

name

The name of the container. If the name isn't specified, the default name "Default" is used. Each container in a pod must have a unique name.

Type: String

Required: No

resources

The type and amount of resources to assign to a container. The supported resources include memory, cpu, and nvidia.com/gpu. For more information, see <a href="Resource">Resource</a> <a href="mailto:management for pods and containers">management for pods and containers</a> in the Kubernetes documentation.

Type: EksContainerResourceRequirements object

Required: No

limits

The type and quantity of the resources to reserve for the container. The values vary based on the name that's specified. Resources can be requested using either the limits or the requests objects.

memory

The memory hard limit (in MiB) for the container, using whole integers, with a "Mi" suffix. If your container attempts to exceed the memory specified, the container is terminated. You must specify at least 4 MiB of memory for a job. memory can be specified in limits, requests, or both. If memory is specified in both places, then the value that's specified in limits must be equal to the value that's specified in requests.



# Note

To maximize your resource utilization, provide your jobs with as much memory as possible for the specific instance type that you are using. To learn how, see Compute resource memory management.

cpu

The number of CPUs that's reserved for the container. Values must be an even multiple of 0.25. cpu can be specified in limits, requests, or both. If cpu is specified in both places, then the value that's specified in limits must be at least as large as the value that's specified in requests.

nvidia.com/gpu

The number of GPUs that's reserved for the container. Values must be a whole integer. memory can be specified in limits, requests, or both. If memory is specified in both places, then the value that's specified in limits must be equal to the value that's specified in requests.

Type: String to string map

Value Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

## requests

The type and quantity of the resources to request for the container. The values vary based on the name that's specified. Resources can be requested by using either the limits or the requests objects.

#### memory

The memory hard limit (in MiB) for the container, using whole integers, with a "Mi" suffix. If your container attempts to exceed the memory specified, the container is terminated. You must specify at least 4 MiB of memory for a job. memory can be specified in limits, requests, or both. If memory is specified in both, then the value that's specified in limits must be equal to the value that's specified in requests.



# Note

If you're trying to maximize your resource utilization by providing your jobs as much memory as possible for a particular instance type, see Compute resource memory management.

#### cpu

The number of CPUs that are reserved for the container. Values must be an even multiple of 0.25. cpu can be specified in limits, requests, or both. If cpu is specified in both, then the value that's specified in limits must be at least as large as the value that's specified in requests.

# nvidia.com/gpu

The number of GPUs that are reserved for the container. Values must be a whole integer. nvidia.com/qpu can be specified in limits, requests, or both. If nvidia.com/gpu is specified in both, then the value that's specified in limits must be equal to the value that's specified in requests.

Type: String to string map

Value Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

securityContext

The security context for a job. For more information, see <u>Configure a security context for</u> a pod or container in the *Kubernetes documentation*.

Type: EksContainerSecurityContext object

Required: No privileged

When this parameter is true, the container is given elevated permissions on the host container instance. The level of permissions is similar to the root user permissions. The default value is false. This parameter maps to privileged policy in the Privileged pod security policies in the Kubernetes documentation.

Type: Boolean

Required: No

readOnlyRootFilesystem

When this parameter is true, the container is given read-only access to its root file system. The default value is false. This parameter maps to ReadOnlyRootFilesystem policy in the Volumes and file systems pod security policies in the Kubernetes documentation.

Type: Boolean

Required: No

runAsGroup

When this parameter is specified, the container is run as the specified group ID (gid). If this parameter isn't specified, the default is the group that's specified in the image metadata. This parameter maps to RunAsGroup and MustRunAs policy in the <u>Users</u> and groups pod security policies in the *Kubernetes documentation*.

Type: Long

Required: No

#### runAsNonRoot

When this parameter is specified, the container is run as a user with a uid other than 0. If this parameter isn't specified, so such rule is enforced. This parameter maps to RunAsUser and MustRunAsNonRoot policy in the <u>Users and groups pod security</u> policies in the *Kubernetes documentation*.

Type: Long

Required: No

runAsUser

When this parameter is specified, the container is run as the specified user ID (uid). If this parameter isn't specified, the default is the user that's specified in the image metadata. This parameter maps to RunAsUser and MustRanAs policy in the <u>Users</u> and groups pod security policies in the *Kubernetes documentation*.

Type: Long

Required: No

volumeMounts

The volume mounts for a container for an Amazon EKS job. For more information about volumes and volume mounts in Kubernetes, see <u>Volumes</u> in the *Kubernetes documentation*.

Type: Array of EksContainerVolumeMount objects

Required: No mountPath

The path on the container where the volume is mounted.

Type: String

Required: No

name

The name the volume mount. This must match the name of one of the volumes in the pod.

Type: String

Required: No

readOnly

If this value is true, the container has read-only access to the volume. Otherwise, the container can write to the volume. The default value is false.

Type: Boolean

Required: No

dnsPolicy

The DNS policy for the pod. The default value is ClusterFirst. If the hostNetwork parameter is not specified, the default is ClusterFirstWithHostNet. ClusterFirst indicates that any DNS query that does not match the configured cluster domain suffix is forwarded to the upstream nameserver inherited from the node. If no value was specified for dnsPolicy in the <a href="RegisterJobDefinition">RegisterJobDefinition</a> API operation, then no value is returned for dnsPolicy by either of <a href="DescribeJobDefinitions">DescribeJobDefinitions</a> or <a href="DescribeJobDefinitions">DescribeJobS</a> API operations. The pod spec setting will contain either ClusterFirst or ClusterFirstWithHostNet, depending on the value of the hostNetwork parameter. For more information, see <a href="Pod's DNS policy">Pod's DNS policy</a> in the <a href="Kubernetes documentation">Kubernetes documentation</a>.

Valid values: Default | ClusterFirst | ClusterFirstWithHostNet

Type: String

Required: No

hostNetwork

Indicates if the pod uses the hosts' network IP address. The default value is true. Setting this to false enables the Kubernetes pod networking model. Most Amazon Batch workloads are egress-only and don't require the overhead of IP allocation for each pod for incoming connections. For more information, see <a href="Host namespaces">Host namespaces</a> and <a href="Pod networking">Pod networking</a> in the Kubernetes documentation.

Type: Boolean

Required: No

serviceAccountName

The name of the service account that's used to run the pod. For more information, see Kubernetes service accounts and Configure a Kubernetes service account to assume an IAM

<u>role</u> in the *Amazon EKS User Guide* and <u>Configure service accounts for pods</u> in the *Kubernetes documentation*.

Type: String

Required: No

volumes

Specifies the volumes for a job definition that uses Amazon EKS resources.

Type: Array of EksVolume objects

Required: No

emptyDir

Specifies the configuration of a Kubernetes emptyDir volume. An emptyDir volume is first created when a pod is assigned to a node. It exists as long as that pod runs on that node. The emptyDir volume is initially empty. All containers in the pod can read and write the files in the emptyDir volume. However, the emptyDir volume can be mounted at the same or different paths in each container. When a pod is removed from a node for any reason, the data in the emptyDir is deleted permanently. For more information, see emptyDir in the *Kubernetes documentation*.

Type: EksEmptyDir object

Required: No

medium

The medium to store the volume. The default value is an empty string, which uses the storage of the node.

,,,,

(**Default**) Use the disk storage of the node.

"Memory"

Use the tmpfs volume that's backed by the RAM of the node. Contents of the volume are lost when the node reboots, and any storage on the volume counts against the container's memory limit.

Type: String

Required: No

sizeLimit

The maximum size of the volume. By default, there's no maximum size defined.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

hostPath

Specifies the configuration of a Kubernetes hostPath volume. A hostPath volume mounts an existing file or directory from the host node's filesystem into your pod. For more information, see hostPath in the *Kubernetes documentation*.

Type: EksHostPath object

Required: No

path

The path of the file or directory on the host to mount into containers on the pod.

Type: String

Required: No

name

The name of the volume. The name must be allowed as a DNS subdomain name. For more information, see DNS subdomain names in the *Kubernetes documentation*.

Type: String

Required: Yes

secret

Specifies the configuration of a Kubernetes secret volume. For more information, see secret in the *Kubernetes documentation*.

Type: EksSecret object

Required: No

## optional

Specifies whether the secret or the secret's keys must be defined.

Type: Boolean

Required: No

secretName

The name of the secret. The name must be allowed as a DNS subdomain name. For more information, see DNS subdomain names in the Kubernetes documentation.

Type: String

Required: Yes

# **Platform capabilities**

platformCapabilities

The platform capabilities that's required by the job definition. If no value is specified, it defaults to EC2. For jobs that run on Fargate resources, FARGATE is specified.



## Note

If the job runs on Amazon EKS resources, then you must not specify platformCapabilities.

Type: String

Valid values: EC2 | FARGATE

Required: No

# Propagate tags

propagateTags

Specifies whether to propagate the tags from the job or job definition to the corresponding Amazon ECS task. If no value is specified, the tags aren't propagated. Tags can only be

propagated to the tasks when the task is created. For tags with the same name, job tags are given priority over job definitions tags. If the total number of combined tags from the job and job definition is over 50, the job's moved to the FAILED state.



#### Note

If the job runs on Amazon EKS resources, then you must not specify propagate Tags.

Type: Boolean

Required: No

# **Node properties**

# nodeProperties

When you register a multi-node parallel job definition, you must specify a list of node properties. These node properties define the number of nodes to use in your job, the main node index, and the different node ranges to use. If the job runs on Fargate resources, then you can't specify nodeProperties. Instead, use containerProperties. The following node properties are allowed in a job definition. For more information, see Multi-node parallel jobs.



## Note

If the job runs on Amazon EKS resources, then you must not specify nodeProperties.

Type: NodeProperties object

Required: No

mainNode

Specifies the node index for the main node of a multi-node parallel job. This node index value must be smaller than the number of nodes.

Type: Integer

Required: Yes

#### numNodes

The number of nodes that are associated with a multi-node parallel job.

Type: Integer

Required: Yes

nodeRangeProperties

A list of node ranges and their properties that are associated with a multi-node parallel job.



# Note

A node group is an identical group of job nodes that all share the same container properties. You can use Amazon Batch to specify up to five distinct node groups for each job.

Type: Array of NodeRangeProperty objects

Required: Yes targetNodes

> The range of nodes, using node index values. A range of 0:3 indicates nodes with index values of 0 through 3. If the starting range value is omitted (:n), then 0 is used to start the range. If the ending range value is omitted (n:), then the highest possible node index is used to end the range. Your accumulative node ranges must account for all nodes (0:n). You can nest node ranges, for example 0:10 and 4:5. For this case, the 4:5 range properties override the 0:10 properties.

Type: String

Required: No

container

The container details for the node range. For more information, see Container properties.

Type: ContainerProperties object

Required: No

Reference: Job definition parameters for ContainerProperties

# **Retry strategy**

# retryStrategy

When you register a job definition, you can optionally specify a retry strategy to use for failed jobs that are submitted with this job definition. Any retry strategy that's specified during a <u>SubmitJob</u> operation overrides the retry strategy defined here. By default, each job is attempted one time. If you specify more than one attempt, the job is retried if it fails. Examples of a fail attempt include the job returns a non-zero exit code or the container instance is terminated. For more information, see <u>Automated job retries</u>.

Type: RetryStrategy object

Required: No

attempts

The number of times to move a job to the RUNNABLE status. You can specify between 1 and 10 attempts. If attempts is greater than one, the job is retried that many times if it fails, until it has moved to RUNNABLE.

```
"attempts": integer
```

Type: Integer

Required: No

evaluateOnExit

Array of up to 5 objects that specify conditions under which the job is retried or failed. If this parameter is specified, then the attempts parameter must also be specified. If evaluateOnExit is specified but none of the entries match, then the job is retried.

Type: Array of EvaluateOnExit objects

Required: No

action

Specifies the action to take if all of the specified conditions (onStatusReason, onReason, and onExitCode) are met. The values aren't case sensitive.

Type: String

Required: Yes

Valid values: RETRY | EXIT

onExitCode

Contains a glob pattern to match against the decimal representation of the ExitCode that's returned for a job. The pattern can be up to 512 characters in length. It can contain only numbers. It cannot contain letters or special characters. It can optionally end with an asterisk (\*) so that only the start of the string needs to be an exact match.

Type: String

Required: No

onReason

Contains a glob pattern to match against the Reason that's returned for a job. The pattern can be up to 512 characters in length. It can contain letters, numbers, periods (.), colons (:), and white space (spaces, tabs). It can optionally end with an asterisk (\*) so that only the start of the string needs to be an exact match.

Type: String

Required: No

onStatusReason

Contains a glob pattern to match against the StatusReason that's returned for a job. The pattern can be up to 512 characters in length. It can contain letters, numbers, periods (.), colons (:), and white space (spaces, tabs). It can optionally end with an asterisk (\*) so that only the start of the string needs to be an exact match.

Type: String

Required: No

### Scheduling priority

schedulingPriority

The scheduling priority for jobs that are submitted with this job definition. This only affects jobs in job queues with a fair-share policy. Jobs with a higher scheduling priority are scheduled before jobs with a lower scheduling priority.

The minimum supported value is 0 and the maximum supported value is 9999.

Type: Integer

Required: No

#### **Tags**

tags

Key-value pair tags to associate with the job definition. For more information, see <u>Tag your</u> Amazon Batch resources.

Type: String to string map

Required: No

#### **Timeout**

timeout

You can configure a timeout duration for your jobs so that if a job runs longer than that, Amazon Batch terminates the job. For more information, see <u>Job timeouts</u>. If a job is terminated because of a timeout, it isn't retried. Any timeout configuration that's specified during a <u>SubmitJob</u> operation overrides the timeout configuration defined here. For more information, see <u>Job timeouts</u>.

Type: JobTimeout object

Required: No

#### attemptDurationSeconds

The time duration in seconds (measured from the job attempt's startedAt timestamp) after Amazon Batch terminates unfinished jobs. The minimum value for the timeout is 60 seconds.

For array jobs, the timeout applies to the child jobs, not to the parent array job.

For multi-node parallel (MNP) jobs, the timeout applies to the whole job, not to the individual nodes.

Type: Integer

Required: No

### Create job definitions using EcsProperties

With Amazon Batch job definitions using <u>EcsProperties</u>, you can model hardware, sensors, 3D environments and other simulations in separate containers. You can use this feature to logically organize your workload components, and separate them from the main application. This feature can be used with Amazon Batch on Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), and Amazon Fargate.

### ContainerProperties versus EcsProperties job definitions

You can choose to use <u>ContainerProperties</u> or <u>EcsProperties</u> job definitions as your use case dictates. At a high-level, running Amazon Batch jobs with EcsProperties is similar to running jobs with a ContainerProperties.

The legacy job definition structure, using ContainerProperties, remains supported. If you currently have workflows using this structure, you can continue to run them.

The main difference is that there is a new object added to the job definition to accommodate EcsProperties-based definitions.

For example, a job definition that uses ContainerProperties on Amazon ECS and Fargate has the following structure:

```
{
    "containerProperties": {
    ...
```

A job definition that uses EcsProperties on Amazon ECS and Fargate has the following structure:

### **General changes to the Amazon Batch APIs**

The following further outlines some of the key differences when using the EcsProperties and the EcsProperties API data types:

- Many of the parameters that are used within ContainerProperties appear within TaskContainerProperties. Some examples include, command, image, privileged, secrets, and users. They can all be found within TaskContainerProperties.
- Some of the TaskContainerProperties parameters don't have functional equivalents in the legacy structure. Some examples include, dependsOn, essential, name, ipcMode, and pidMode. For more information, see EcsTaskDetails and TaskContainerProperties.

As well, some ContainerProperties parameters don't have equivalents, or application, in the EcsProperties structure. In <u>taskProperties</u>, container has been replaced with containers so that the new object can accept up to ten elements. For more information see RegisterJobDefinition:containerProperties and EcsTaskProperties:containers.

 taskRoleArn is functionally equivalent to jobRoleArn. For more information see EcsTaskProperties:taskRoleArn and ContainerProperties:jobRoleArn.

- You can include from one (1) to ten (10) containers in the EcsProperties structure. For more information see EcsTaskProperties:containers.
- The taskProperties and instanceTypes objects are arrays, but currently accept only one element. For example, EcsProperties:taskProperties and NodeRangeProperty:instanceTypes.

### Multi-container job definitions for Amazon ECS

To accommodate the multi-container structure for Amazon ECS, some of the API data types are different. For example,

- ecsProperties is the same level as containerProperties in the single-container definition. For more information, see EcsProperties in the Amazon Batch API Reference Guide.
- taskProperties contains the properties defined for the Amazon ECS task. For more information, see EcsProperties in the Amazon Batch API Reference Guide.
- containers includes similar information to containerProperties in the single-container definition. The main difference is that containers allows you to define up to ten containers. For more information, see ECSTaskProperties:containers in the Amazon Batch API Reference Guide.
- essential parameter indicates how the container affects the job. All essential containers must complete successfully (exit as 0) in order for the job to progress. If a container that is marked as essential fails (exits as non-0), then the job fails.
  - The default value is true and at least one container must be marked as essential. For more information, see essential in the Amazon Batch API Reference Guide.
- With the depends 0n parameter, you can define a list of container dependencies. For more information, see depends0n in the Amazon Batch API Reference Guide.



#### (i) Note

The complexity of the depends 0n list and the associated container runtime can affect the start time for your job. If the dependencies take a long time to run, the job will remain in a STARTING state until they complete.

For more information about the ecsProperties and structure, see RegisterJobDefinition request syntax for ecsProperties.

### Multi-container job definitions for Amazon EKS

To accommodate the multi-container structure for Amazon EKS, some of the API data types are different. For example,

- name is a unique identifier for the container. This object isn't required for a single container, but is required when defining multiple containers in a pod. When name isn't defined for single containers, the default name, default, is applied.
- initContainers are defined within the eksPodProperties data type. They run before application containers, always runs to completion, and must complete successfully before the next container starts.

These containers are registered with the Amazon EKS Connector agent and persists the registration information in the Amazon Elastic Kubernetes Service backend data store. The initContainers object can accept up to ten (10) elements. For more information, see Init Containers in the Kubernetes documentation.

#### Note

The initContainers object can affect the start time for your job. If the initContainers take a long time to run, the job will remain in a STARTING state until they complete.

- shareProcessNamespace indicates if the containers in the pod can share the same process namespace. The default values is false. Setting this to true to enable containers see and signal processes in other containers that located in the same pod.
- Every container has importance. All containers must complete successfully (exit as 0) for the job to succeed. If one container fails (exits as other than 0), then the job fails.

For more information about the eksProperties and structure, see RegisterJobDefinition request syntax for eksProperties.

### Reference: Amazon Batch job scenarios using EcsProperties

To illustrate how Amazon Batch job definitions that use EcsProperties can be structured based on your needs, this topic presents the following <a href="RegisterJobDefinition">RegisterJobDefinition</a> payloads. You can copy these examples into a file, customize them to your needs, and then use the Amazon Command Line Interface (Amazon CLI) to call RegisterJobDefinition.

### **Amazon Batch job for Amazon ECS on Amazon EC2**

The following is an example of a Amazon Batch job for Amazon Elastic Container Service on Amazon Elastic Compute Cloud:

```
{
    "jobDefinitionName": "multicontainer-ecs-ec2",
    "type": "container",
    "ecsProperties": {
        "taskProperties": [
            "containers": [
                 "name": "c1",
                "essential": false,
                 "command": [
                   "echo",
                   "hello world"
                ],
                "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
                 "resourceRequirements": [
                  {
                     "type": "VCPU",
                     "value": "2"
                   },
                     "type": "MEMORY",
                     "value": "4096"
                   }
                ]
              },
                "name": "c2",
                "essential": false,
                 "command": [
                   "echo",
```

```
"hello world"
        ],
        "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
        "resourceRequirements": [
          {
            "type": "VCPU",
            "value": "2"
          },
            "type": "MEMORY",
            "value": "4096"
          }
        ]
      },
        "name": "c3",
        "essential": true,
        "command": [
          "echo",
          "hello world"
        ],
        "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
        "firelensConfiguration": {
          "type": "fluentbit",
          "options": {
            "enable-ecs-log-metadata": "true"
          }
         },
        "resourceRequirements": [
          {
            "type": "VCPU",
            "value": "6"
          },
            "type": "MEMORY",
            "value": "12288"
          }
        ]
      }
    ]
 }
]
```

}

### **Amazon Batch job for Amazon ECS on Fargate**

The following is an example of a Amazon Batch job for Amazon Elastic Container Service on Amazon Fargate:

```
{
    "jobDefinitionName": "multicontainer-ecs-fargate",
    "type": "container",
    "platformCapabilities": [
        "FARGATE"
    ],
    "ecsProperties": {
        "taskProperties": [
          {
            "containers": [
                "name": "c1",
                "command": [
                  "echo",
                  "hello world"
                ],
                "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
                "resourceRequirements": [
                  {
                     "type": "VCPU",
                    "value": "2"
                  },
                    "type": "MEMORY",
                    "value": "4096"
                ]
              },
                "name": "c2",
                "essential": true,
                "command": [
                  "echo",
                  "hello world"
                ],
                "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
```

```
"resourceRequirements": [
                   {
                     "type": "VCPU",
                     "value": "6"
                   },
                   {
                     "type": "MEMORY",
                     "value": "12288"
                ]
              }
            ],
            "executionRoleArn": "arn:aws:iam::1112223333:role/ecsTaskExecutionRole"
          }
        ]
  }
}
```

### **Amazon Batch job for Amazon EKS**

The following is an example of a Amazon Batch job for Amazon Elastic Kubernetes Service:

```
{
  "jobDefinitionName": "multicontainer-eks",
  "type": "container",
  "eksProperties": {
    "podProperties": {
      "shareProcessNamespace": true,
      "initContainers": [
        {
          "name": "init-container",
          "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
          "command": [
            "echo"
          ],
          "args": [
            "hello world"
          ],
          "resources": {
            "requests": {
              "cpu": "1",
              "memory": "512Mi"
            }
```

```
}
  },
    "name": "init-container-2",
    "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
    "command": [
      "echo",
      "my second init container"
    ],
    "resources": {
      "requests": {
        "cpu": "1",
        "memory": "512Mi"
      }
    }
  }
],
"containers": [
 {
    "name": "c1",
    "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
    "command": [
      "echo world"
   ],
    "resources": {
      "requests": {
        "cpu": "1",
        "memory": "512Mi"
      }
    }
  },
    "name": "sleep-container",
    "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
    "command": [
      "sleep",
      "20"
    "resources": {
      "requests": {
        "cpu": "1",
        "memory": "512Mi"
      }
    }
```

### MNP Amazon Batch job with multiple containers per node

The following is an example of a multi-node parallel (MNP) Amazon Batch job with multiple containers per node:

```
{
  "jobDefinitionName": "multicontainer-mnp",
  "type": "multinode",
  "nodeProperties": {
    "numNodes": 6,
    "mainNode": 0,
    "nodeRangeProperties": [
      {
        "targetNodes": "0:5",
        "ecsProperties": {
          "taskProperties": [
            {
              "containers": [
                {
                   "name": "range05-c1",
                  "command": [
                     "echo",
                     "hello world"
                  ],
                  "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
                   "resourceRequirements": [
                     {
                       "type": "VCPU",
                       "value": "2"
                     },
                     {
                       "type": "MEMORY",
                       "value": "4096"
                     }
                  ]
                },
```

```
"name": "range05-c2",
                   "command": [
                     "echo",
                     "hello world"
                   ],
                   "image": "public.ecr.aws/amazonlinux/amazonlinux:latest",
                   "resourceRequirements": [
                     {
                       "type": "VCPU",
                       "value": "2"
                     },
                     {
                       "type": "MEMORY",
                       "value": "4096"
                     }
                   ]
                 }
               ]
            }
          ]
        }
      }
    ]
  }
}
```

# Use the awslogs log driver

By default, Amazon Batch enables the awslogs log driver to send log information to CloudWatch Logs. You can use this feature to view different logs from your containers in one convenient location and prevent your container logs from taking up disk space on your container instances. This topic helps you configure the awslogs log driver in your job definitions.



In the Amazon Batch console, you can configure the awslogs log driver in the Logging **configuration** section when you create a job definition.

Use the awslogs log driver 252



#### Note

The type of information that's logged by the containers in your job depends mostly on their ENTRYPOINT command. By default, the logs that are captured show the command output that you normally see in an interactive terminal if you ran the container locally, which are the STDOUT and STDERR I/O streams. The awslogs log driver simply passes these logs from Docker to CloudWatch Logs. For more information about how Docker logs are processed, including alternative ways to capture different file data or streams, see View logs for a container or service in the Docker documentation.

To send system logs from your container instances to CloudWatch Logs, see Using CloudWatch Logs with Amazon Batch. For more information about CloudWatch Logs, see Monitoring Log Files and CloudWatch Logs quotas in the Amazon CloudWatch Logs User Guide.

### awslogs log driver options in the Amazon Batch JobDefiniton data type

The awslogs log driver supports the following options in Amazon Batch job definitions. For more information, see CloudWatch Logs logging driver in the Docker documentation.

awslogs-region

Required: No

Specify the Region where the awslogs log driver should send your Docker logs. By default, the Region that's used is the same one as the one for the job. You can choose to send all of your logs from jobs in different Regions to a single Region in CloudWatch Logs. Doing this allows them to be visible all from one location. Alternatively, you can separate them by Region for more granular approach. However, when you choose this option, make sure that the specified log groups exists in the Region that you specified.

awslogs-group

Required: Optional

With the awslogs-group option, you can specify the log group that the awslogs log driver sends its log streams to. If this isn't specified, aws/batch/job is used.

awslogs-stream-prefix

Required: Optional

With the awslogs-stream-prefix option, you can associate a log stream with the specified prefix, and the Amazon ECS task ID of the Amazon Batch job that the container belongs to. If you specify a prefix with this option, then the log stream takes the following format:

prefix-name/default/ecs-task-id

awslogs-datetime-format

Required: No

This option defines a multiline start pattern in Python strftime format. A log message consists of a line that matches the pattern and any following lines that don't match the pattern. Thus the matched line is the delimiter between log messages.

One example of a use case for using this format is for parsing output such as a stack dump, which might otherwise be logged in multiple entries. The correct pattern allows it to be captured in a single entry.

For more information, see awslogs-datetime-format.

This option always takes precedence if both awslogs-datetime-format and awslogsmultiline-pattern are configured.



#### Note

Multiline logging performs regular expression parsing and matching of all log messages. This may have a negative impact on logging performance.

awslogs-multiline-pattern

Required: No

This option defines a multiline start pattern using a regular expression. A log message consists of a line that matches the pattern and any following lines that don't match the pattern. Thus, the matched line is the delimiter between log messages.

For more information, see awslogs-multiline-pattern in the Docker documentation.

This option is ignored if awslogs-datetime-format is also configured.



#### Note

Multiline logging performs regular expression parsing and matching of all log messages. This might have a negative impact on logging performance.

awslogs-create-group

Required: No

Specify whether you want the log group automatically created. If this option isn't specified, it defaults to false.



#### Marning

This option isn't recommended. We recommend that you create the log group in advance using the CloudWatch Logs CreateLogGroup API action as each job tries to create the log group, increasing the chance that the job fails.



#### Note

The IAM policy for your execution role must include the logs:CreateLogGroup permission before you attempt to use awslogs-create-group.

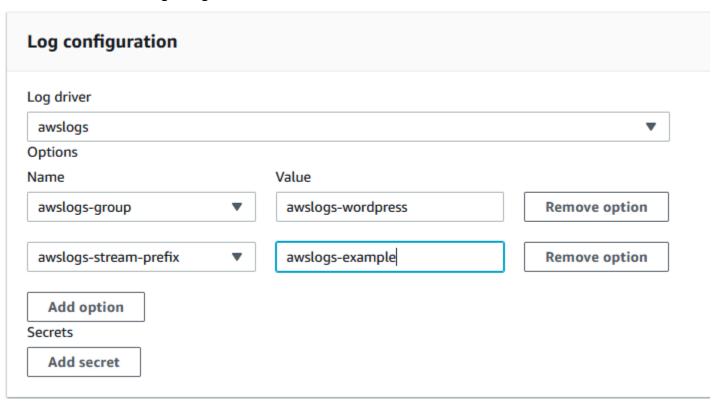
# Specify a log configuration in your job definition

By default, Amazon Batch enables the awslogs log driver. This section describes how to customize the awslogs log configuration for a job. For more information, see Create a single-node job definition.

The following log configuration JSON snippets have a logConfiguration object specified for each job. One is for a WordPress job that sends logs to a log group called awslogs-wordpress, and another is for a MySQL container that sends logs to a log group called awslogs-mysql. Both containers use the awslogs-example log stream prefix.

"logConfiguration": {

In the Amazon Batch console, the log configuration for the wordpress job definition is specified as shown in the following image.



After you have registered a task definition with the awslogs log driver in a job definition log configuration, you can submit a job with that job definition to start sending logs to CloudWatch Logs. For more information, see Tutorial: submit a job.

# Specify sensitive data

With Amazon Batch, you can inject sensitive data into your jobs by storing your sensitive data in either Amazon Secrets Manager secrets or Amazon Systems Manager Parameter Store parameters, and then reference them in your job definition.

Secrets can be exposed to a job in the following ways:

- To inject sensitive data into your containers as environment variables, use the secrets job definition parameter.
- To reference sensitive information in the log configuration of a job, use the secretOptions job definition parameter.

#### **Topics**

- Specify sensitive data with Secrets Manager
- Specify sensitive data with Systems Manager Parameter Store

### **Specify sensitive data with Secrets Manager**

With Amazon Batch, you can inject sensitive data into your jobs by storing your sensitive data in Amazon Secrets Manager secrets and then referencing them in your job definition. Sensitive data stored in Secrets Manager secrets can be exposed to a job as environment variables or as part of the log configuration.

When you inject a secret as an environment variable, you can specify a JSON key or version of a secret to inject. This process helps you control the sensitive data exposed to your job. For more information about secret versioning, see <a href="Key Terms and Concepts for Amazon Secrets Manager">Key Terms and Concepts for Amazon Secrets Manager</a> in the Amazon Secrets Manager User Guide.

### Considerations when you specify sensitive data using Secrets Manager

The following should be considered when using Secrets Manager to specify sensitive data for jobs.

To inject a secret using a specific JSON key or version of a secret, the container instance in your compute environment must have version 1.37.0 or later of the Amazon ECS container agent installed. However, we recommend using the latest container agent version. For information about checking your agent version and updating to the latest version, see <a href="Updating the Amazon Elastic Container Service Developer Guide">Updating the Amazon Elastic Container Service Developer Guide</a>.

Specify sensitive data 257

To inject the full contents of a secret as an environment variable or to inject a secret in a log configuration, your container instance must have version 1.23.0 or later of the container agent.

- Only secrets that store text data, which are secrets created with the SecretString parameter
  of the <u>CreateSecret</u> API, are supported. Secrets that store binary data, which are secrets created
  with the SecretBinary parameter of the <u>CreateSecret</u> API aren't supported.
- When using a job definition that references Secrets Manager secrets to retrieve sensitive data
  for your jobs, if you're also using interface VPC endpoints, you must create the interface VPC
  endpoints for Secrets Manager. For more information, see <u>Using Secrets Manager with VPC</u>
  <u>Endpoints</u> in the *Amazon Secrets Manager User Guide*.
- Sensitive data is injected into your job when the job is initially started. If the secret is
  subsequently updated or rotated, the job doesn't receive the updated value automatically. You
  must launch a new job to force the service to launch a fresh job with the updated secret value.

#### **Required IAM permissions for Amazon Batch secrets**

To use this feature, you must have the execution role and reference it in your job definition. This allows the container agent to pull the necessary Secrets Manager resources. For more information, see Amazon Batch IAM execution role.

To provide access to the Secrets Manager secrets that you create, manually add the following permissions as an inline policy to the execution role. For more information, see <a href="Adding and Removing IAM Policies">Adding and Removing IAM Policies</a> in the IAM User Guide.

- secretsmanager:GetSecretValue—Required if you're referencing a Secrets Manager secret.
- kms:Decrypt—Required only if your secret uses a custom KMS key and not the default key. The ARN for your custom key should be added as a resource.

The following example inline policy adds the required permissions.

**JSON** 

```
{
  "Version": "2012-10-17",
  "Statement": [
     {
        "Effect": "Allow",
```

```
"Action": [
    "secretsmanager:GetSecretValue",
    "kms:Decrypt"
],
    "Resource": [
        "arn:aws-
cn:secretsmanager:<region>:<aws_account_id>:secret:<secret_name>",
        "arn:aws-cn:kms:<region>:<aws_account_id>:key/<key_id>"
]
}
]
}
```

### Injecting sensitive data as an environment variable

You can specify the following Within your job definition:

- The secrets object containing the name of the environment variable to set in the job
- The Amazon Resource Name (ARN) of the Secrets Manager secret
- Additional parameters that contain the sensitive data to present to the job

The following example shows the full syntax that must be specified for the Secrets Manager secret.

```
arn:aws-cn:secretsmanager:region:aws_account_id:secret:secret-name:json-key:version-
stage:version-id
```

The following section describes the additional parameters. These parameters are optional. However, if you don't use them, you must include the colons: to use the default values. Examples are provided below for more context.

```
json-key
```

Specifies the name of the key in a key-value pair with the value that you want to set as the environment variable value. Only values in JSON format are supported. If you don't specify a JSON key, then the full contents of the secret is used.

#### version-stage

Specifies the staging label of the version of a secret that you want to use. If a version staging label is specified, you can't specify a version ID. If no version stage is specified, the default behavior is to retrieve the secret with the AWSCURRENT staging label.

Staging labels are used to keep track of different versions of a secret when they are either updated or rotated. Each version of a secret has one or more staging labels and an ID. For more information, see <a href="Key Terms and Concepts for Amazon Secrets Manager">Key Terms and Concepts for Amazon Secrets Manager</a> in the Amazon Secrets Manager User Guide.

```
version-id
```

Specifies the unique identifier of the version of a secret that you want to use. If a version ID is specified, you can't specify a version staging label. If no version ID is specified, the default behavior is to retrieve the secret with the AWSCURRENT staging label.

Version IDs are used to keep track of different versions of a secret when they are either updated or rotated. Each version of a secret has an ID. For more information, see <a href="Key Terms and Concepts">Key Terms and Concepts for Amazon Secrets Manager</a> in the Amazon Secrets Manager User Guide.

#### **Example container definitions**

The following examples show ways that you can reference Secrets Manager secrets in your container definitions.

#### Example referencing a full secret

The following is a snippet of a task definition showing the format when referencing the full text of a Secrets Manager secret.

```
{
  "containerProperties": [{
      "secrets": [{
            "name": "environment_variable_name",
            "valueFrom": "arn:aws-cn:secretsmanager:region:aws_account_id:secret:secret_name-
AbCdEf"
      }]
  }]
}
```

#### Example referencing a specific key within a secret

The following shows an example output from a <u>get-secret-value</u> command that displays the contents of a secret along with the version staging label and version ID associated with it.

Reference a specific key from the previous output in a container definition by specifying the key name at the end of the ARN.

```
{
  "containerProperties": [{
      "secrets": [{
            "name": "environment_variable_name",
            "valueFrom": "arn:aws-
cn:secretsmanager:region:aws_account_id:secret:appauthexample-AbCdEf:username1::"
      }]
  }]
}
```

#### Example referencing a specific secret version

The following shows an example output from a <u>describe-secret</u> command that displays the unencrypted contents of a secret along with the metadata for all versions of the secret.

```
{
    "ARN": "arn:aws-cn:secretsmanager:region:aws_account_id:secret:appauthexample-
AbCdEf",
    "Name": "appauthexample",
    "Description": "Example of a secret containing application authorization data.",
```

Reference a specific version staging label from the previous output in a container definition by specifying the key name at the end of the ARN.

```
{
   "containerProperties": [{
        "secrets": [{
            "name": "environment_variable_name",
            "valueFrom": "arn:aws-
cn:secretsmanager:region:aws_account_id:secret:appauthexample-AbCdEf::AWSPREVIOUS:"
      }]
   }]
}
```

Reference a specific version ID from the previous output in a container definition by specifying the key name at the end of the ARN.

```
{
  "containerProperties": [{
      "secrets": [{
            "name": "environment_variable_name",
            "valueFrom": "arn:aws-
cn:secretsmanager:region:aws_account_id:secret:appauthexample-AbCdEf::9d4cb84b-ad69-40c0-a0ab-cead36b967e8"
      }]
    }]
}
```

#### Example referencing a specific key and version staging label of a secret

The following shows how to reference both a specific key within a secret and a specific version staging label.

To specify a specific key and version ID, use the following syntax.

```
{
  "containerProperties": [{
      "secrets": [{
            "name": "environment_variable_name",
            "valueFrom": "arn:aws-
cn:secretsmanager:region:aws_account_id:secret:appauthexample-
AbCdEf:username1::9d4cb84b-ad69-40c0-a0ab-cead36b967e8"
      }]
    }]
}
```

### Inject sensitive data in a log configuration

When you specify a logConfiguration within your job definition, you can specify secretOptions with the name of the log driver option to set in the container and the full ARN of the Secrets Manager secret containing the sensitive data to present to the container.

The following is a snippet of a job definition showing the format when referencing an Secrets Manager secret.

```
{
  "containerProperties": [{
    "logConfiguration": [{
```

### Create an Amazon Secrets Manager secret

You can use the Secrets Manager console to create a secret for your sensitive data. For more information, see <a href="Creating a Basic Secret">Creating a Basic Secret</a> in the Amazon Secrets Manager User Guide.

#### To create a basic secret

Use Secrets Manager to create a secret for your sensitive data.

- 1. Open the Secrets Manager console at https://console.amazonaws.cn/secretsmanager/.
- Choose Store a new secret.
- 3. For **Select secret type**, choose **Other type of secrets**.
- 4. Specify the details of your custom secret as **Key** and **Value** pairs. For example, you can specify a key of UserName, and then supply the appropriate user name as its value. Add a second key with the name of Password and the password text as its value. You could also add entries for a database name, server address, or TCP port. You can add as many pairs as you need to store the information you require.

Alternatively, you can choose the **Plaintext** tab and enter the secret value in any way you like.

- 5. Choose the Amazon KMS encryption key that you want to use to encrypt the protected text in the secret. If you don't choose one, Secrets Manager checks to see if there's a default key for the account, and uses it if it exists. If a default key doesn't exist, Secrets Manager creates one for you automatically. You can also choose **Add new key** to create a custom KMS key specifically for this secret. To create your own KMS key, you must have permissions to create KMS keys in your account.
- Choose Next.

- 7. For **Secret name**, type an optional path and name, such as **production/ MyAwesomeAppSecret** or **development/TestSecret**, and choose **Next**. You can optionally add a description to help you remember the purpose of this secret later.
  - The secret name must be ASCII letters, digits, or any of the following characters: /\_+=.@-
- 8. (Optional) At this point, you can configure rotation for your secret. For this procedure, leave it at **Disable automatic rotation** and choose **Next**.
  - For information about how to configure rotation on new or existing secrets, see <u>Rotating Your Amazon Secrets Manager Secrets</u>.
- 9. Review your settings, and then choose **Store secret** to save everything you entered as a new secret in Secrets Manager.

### Specify sensitive data with Systems Manager Parameter Store

With Amazon Batch, you can inject sensitive data into your containers by storing your sensitive data in Amazon Systems Manager Parameter Store parameters and then referencing them in your container definition.

#### **Topics**

- Considerations when you specify sensitive data using Systems Manager Parameter Store
- Required IAM permissions for Amazon Batch secrets
- Inject sensitive data as an environment variable
- Inject sensitive data in a log configuration
- Create an Amazon Systems Manager Parameter Store parameter

# Considerations when you specify sensitive data using Systems Manager Parameter Store

The following should be considered when specifying sensitive data for containers using Systems Manager Parameter Store parameters.

This feature requires that your container instance have version 1.23.0 or later of the container
agent. However, we recommend using the latest container agent version. For information about
checking your agent version and updating to the latest version, see <u>Updating the Amazon ECS</u>
container agent in the Amazon Elastic Container Service Developer Guide.

• Sensitive data is injected into the container for your job when the container is initially started. If the secret or Parameter Store parameter is subsequently updated or rotated, the container doesn't receive the updated value automatically. You must launch a new job to force the launch of a fresh job with updated secrets.

### **Required IAM permissions for Amazon Batch secrets**

To use this feature, you must have the execution role and reference it in your job definition. This allows the Amazon ECS container agent to pull the necessary Amazon Systems Manager resources. For more information, see Amazon Batch IAM execution role.

To provide access to the Amazon Systems Manager Parameter Store parameters that you create, manually add the following permissions as an inline policy to the execution role. For more information, see Adding and Removing IAM Policies in the IAM User Guide.

- ssm:GetParameters—Required if you're referencing a Systems Manager Parameter Store parameter in a task definition.
- secretsmanager: GetSecretValue—Required if you're referencing a Secrets Manager secret either directly or if your Systems Manager Parameter Store parameter is referencing a Secrets Manager secret in a task definition.
- kms:Decrypt—Required only if your secret uses a custom KMS key and not the default key. The ARN for your custom key should be added as a resource.

The following example inline policy adds the required permissions:

**JSON** 

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
       "Effect": "Allow",
       "Action": [
            "ssm:GetParameters",
            "secretsmanager:GetSecretValue",
            "kms:Decrypt"
        ],
        "Resource": [
```

#### Inject sensitive data as an environment variable

Within your container definition, specify secrets with the name of the environment variable to set in the container and the full ARN of the Systems Manager Parameter Store parameter containing the sensitive data to present to the container.

The following is a snippet of a task definition showing the format when referencing an Systems Manager Parameter Store parameter. If the Systems Manager Parameter Store parameter exists in the same Region as the task that you're launching, then you can use either the full ARN or name of the parameter. If the parameter exists in a different Region, then the full ARN must be specified.

```
{
  "containerProperties": [{
     "secrets": [{
        "name": "environment_variable_name",
        "valueFrom": "arn:aws-cn:ssm:region:aws_account_id:parameter/parameter_name"
     }]
  }]
}
```

### Inject sensitive data in a log configuration

Within your container definition, when specifying a logConfiguration you can specify secretOptions with the name of the log driver option to set in the container and the full ARN of the Systems Manager Parameter Store parameter containing the sensitive data to present to the container.

#### Important

If the Systems Manager Parameter Store parameter exists in the same Region as the task you're launching, then you can use either the full ARN or name of the parameter. If the parameter exists in a different Region, then the full ARN must be specified.

The following is a snippet of a task definition showing the format when referencing an Systems Manager Parameter Store parameter.

```
{
  "containerProperties": [{
    "logConfiguration": [{
      "logDriver": "fluentd",
      "options": {
        "tag": "fluentd demo"
      },
      "secretOptions": [{
        "name": "fluentd-address",
        "valueFrom": "arn:aws-cn:ssm:region:aws_account_id:parameter/parameter_name"
      }]
    }]
  }]
}
```

### Create an Amazon Systems Manager Parameter Store parameter

You can use the Amazon Systems Manager console to create a Systems Manager Parameter Store parameter for your sensitive data. For more information, see Walkthrough: Create and Use a Parameter in a Command (Console) in the Amazon Systems Manager User Guide.

#### To create a Parameter Store parameter

- Open the Amazon Systems Manager console at https://console.amazonaws.cn/systems-1. manager/.
- 2. In the navigation pane, choose **Parameter Store**, **Create parameter**.
- 3. For **Name**, type a hierarchy and a parameter name. For example, type test/ database\_password.
- For **Description**, type an optional description.

5. For Type, choose String, StringList, or SecureString.



If you choose SecureString, the KMS Key ID field appears. If you don't provide
 a KMS key ID, a KMS key ARN, an alias name, or an alias ARN, then the system
 uses alias/aws/ssm. This is the default KMS key for Systems Manager. To avoid
 using this key, choose a custom key. For more information, see <u>Use Secure String</u>
 Parameters in the Amazon Systems Manager User Guide.

 When you create a secure string parameter in the console by using the key-id parameter with either a custom KMS key alias name or an alias ARN, you must specify the prefix alias/before the alias. The following is an ARN example:

```
arn:aws-cn:kms:us-east-2:123456789012:alias/MyAliasName
```

The following is an alias name example:

```
alias/MyAliasName
```

- 6. For **Value**, type a value. For example, MyFirstParameter. If you chose **SecureString**, the value is masked exactly as you entered it.
- 7. Choose **Create parameter**.

# Private registry authentication for jobs

Private registry authentication for jobs using Amazon Secrets Manager enables you to store your credentials securely and then reference them in your job definition. This provides a way to reference container images that exist in private registries outside of Amazon that require authentication in your job definitions. This feature is supported by jobs hosted on Amazon EC2 instances and Fargate.

#### Important

If your job definition references an image that's stored in Amazon ECR, this topic doesn't apply. For more information, see Using Amazon ECR Images with Amazon ECS in the Amazon Elastic Container Registry User Guide.

For jobs hosted on Amazon EC2 instances, this feature requires version 1.19.0 or later of the container agent. However, we recommend using the latest container agent version. For information about how to check your agent version and update to the latest version, see Updating the Amazon ECS container agent in the Amazon Elastic Container Service Developer Guide.

For jobs hosted on Fargate, this feature requires platform version 1.2.0 or later. For information, see Amazon Fargate Linux platform versions in the Amazon Elastic Container Service Developer Guide.

Within your container definition, specify the repositoryCredentials object with the details of the secret that you created. The secret you reference can be from a different Amazon Web Services Region or a different account than the job using it.



When using the Amazon Batch API, Amazon CLI, or Amazon SDK, if the secret exists in the same Amazon Web Services Region as the job that you're launching then you can use either the full ARN or name of the secret. If the secret exists in a different account, the full ARN of the secret must be specified. When using the Amazon Web Services Management Console, the full ARN of the secret must be specified always.

The following is a snippet of a job definition that shows the required parameters:

```
"containerProperties": [
    "image": "private-repo/private-image",
    "repositoryCredentials": {
      "credentialsParameter": "arn:aws-
cn:secretsmanager:region:123456789012:secret:secret_name"
    }
 }
```

]

### Required IAM permissions for private registry authentication

The execution role is required to use this feature. This allows the container agent to pull the container image. For more information, see Amazon Batch IAM execution role.

To provide access to the secrets that you create, add the following permissions as an inline policy to the execution role. For more information, see Adding and Removing IAM Policies.

- secretsmanager:GetSecretValue
- kms:Decrypt—Required only if your key uses a custom KMS key and not the default key. The Amazon Resource Name (ARN) for your custom key must be added as a resource.

The following is an example inline policy that adds the permissions.

**JSON** 

### Tutorial: Create a secret for private registry authentication

Complete the following steps to create a secret for your private registry credentials with Amazon Secrets Manager.

#### Create a basic secret

- Open the Amazon Secrets Manager console at <a href="https://console.amazonaws.cn/">https://console.amazonaws.cn/</a> secretsmanager/.
- Choose Store a new secret.
- 3. For **Select secret type**, choose **Other type of secrets**.
- 4. Select **Plaintext** and enter your private registry credentials using the following format:

```
{
   "username" : "privateRegistryUsername",
   "password" : "privateRegistryPassword"
}
```

- Choose Next.
- 6. For **Secret name**, enter an optional path and name, such as **production/ MyAwesomeAppSecret** or **development/TestSecret**, and choose **Next**. You can optionally add a description to help you remember the purpose of this secret later.
  - The secret name must be ASCII letters, digits, or any of the following characters: /\_+=.@-.
- 7. (Optional) At this point, you can configure rotation for your secret. For this procedure, leave it at **Disable automatic rotation** and choose **Next**.
  - For instructions on how to configure rotation on new or existing secrets, see Rotating Your Amazon Secrets Manager Secrets.
- 8. Review your settings, and then choose **Store secret** to save everything that you entered as a new secret in Secrets Manager.

Register a job definition and under **Private registry**, turn on **Private registry authentication**. Then, in **Secrets Manager ARN or name**, enter the Amazon Resource Name (ARN) of the secret. For more information, see Required IAM permissions for private registry authentication.

### **Amazon EFS volumes**

Amazon Elastic File System (Amazon EFS) provides simple, scalable file storage for use with your Amazon Batch jobs. With Amazon EFS, storage capacity is elastic. It scales automatically as you add and remove files. Your applications can have the storage they need, when they need it.

Amazon EFS volumes 272

You can use Amazon EFS file systems with Amazon Batch to export file system data across your fleet of container instances. That way, your jobs have access to the same persistent storage. However, you must configure your container instance AMI to mount the Amazon EFS file system before the Docker daemon starts. Also, your job definitions must reference volume mounts on the container instance to use the file system. The following sections help you get started using Amazon EFS with Amazon Batch.

#### **Amazon EFS volume considerations**

The following should be considered when using Amazon EFS volumes:

For jobs using EC2 resources, Amazon EFS file system support was added as a public preview with Amazon ECS optimized AMI version 20191212 with container agent version 1.35.0. However, Amazon EFS file system support entered general availability with Amazon ECS optimized AMI version 20200319 with container agent version 1.38.0, which contained the Amazon EFS access point and IAM authorization features. We recommend that you use Amazon ECS optimized AMI version 20200319 or later to take advantage of these features. For more information, see <a href="Amazon ECS optimized AMI versions">Amazon ECS optimized AMI versions</a> in the Amazon Elastic Container Service Developer Guide.

### Note

If you create your own AMI, you must use container agent 1.38.0 or later, ecs-init version 1.38.0-1 or later, and run the following commands on your Amazon EC2 instance. This is all to enable the Amazon ECS volume plugin. The commands are dependent on whether you're using Amazon Linux 2 or Amazon Linux as your base image.

Amazon Linux 2

```
$ yum install amazon-efs-utils
systemctl enable --now amazon-ecs-volume-plugin
```

Amazon Linux

```
$ yum install amazon-efs-utils
sudo shutdown -r now
```

• For jobs using Fargate resources, Amazon EFS file system support was added when using platform version 1.4.0 or later. For more information, see Amazon Fargate platform versions in the Amazon Elastic Container Service Developer Guide.

• When specifying Amazon EFS volumes in jobs using Fargate resources, Fargate creates a supervisor container that is responsible for managing the Amazon EFS volume. The supervisor container uses a small amount of the job's memory. The supervisor container is visible when querying the task metadata version 4 endpoint. For more information, see Task metadata endpoint version 4 in the Amazon Elastic Container Service User Guide for Amazon Fargate.

### **Use Amazon EFS access points**

Amazon EFS access points are application-specific entry points into an EFS file system that help you to manage application access to shared datasets. For more information about Amazon EFS access points and how to control access to them, see Working with Amazon EFS Access Points in the Amazon Elastic File System User Guide.

Access points can enforce a user identity, including the user's POSIX groups, for all file system requests that are made through the access point. Access points can also enforce a different root directory for the file system so that clients can only access data in the specified directory or its subdirectories.



#### Note

When creating an EFS access point, you specify a path on the file system to serve as the root directory. When you reference the EFS file system with an access point ID in your Amazon Batch job definition, the root directory must either be omitted or set to / This enforces the path that's set on the EFS access point.

You can use an Amazon Batch job IAM role to enforce that specific applications use a specific access point. By combining IAM policies with access points, you can easily provide secure access to specific datasets for your applications. This feature uses Amazon ECS IAM roles for task functionality. For more information, see IAM Roles for Tasks in the Amazon Elastic Container Service Developer Guide.

# Specify an Amazon EFS file system in your job definition

To use Amazon EFS file system volumes for your containers, you must specify the volume and mount point configurations in your job definition. The following job definition JSON snippet shows the syntax for the volumes and mountPoints objects for a container:

```
{
    "containerProperties": [
        {
            "image": "amazonlinux:2",
            "command": [
                "ls",
                "-la",
                "/mount/efs"
            ],
            "mountPoints": [
                {
                     "sourceVolume": "myEfsVolume",
                     "containerPath": "/mount/efs",
                     "readOnly": true
                }
            ],
            "volumes": [
                     "name": "myEfsVolume",
                     "efsVolumeConfiguration": {
                         "fileSystemId": "fs-12345678",
                         "rootDirectory": "/path/to/my/data",
                         "transitEncryption": "ENABLED",
                         "transitEncryptionPort": integer,
                         "authorizationConfig": {
                             "accessPointId": "fsap-1234567890abcdef1",
                             "iam": "ENABLED"
                         }
                     }
                }
            ]
        }
    ]
}
```

#### efsVolumeConfiguration

Type: Object

Required: No

This parameter is specified when using Amazon EFS volumes.

fileSystemId

Type: String

Required: Yes

The Amazon EFS file system ID to use.

rootDirectory

Type: String

Required: No

The directory within the Amazon EFS file system to mount as the root directory inside the host. If this parameter is omitted, the root of the Amazon EFS volume is used. Specifying / has the same effect as omitting this parameter. It can be up to 4,096 characters in length.

#### Important

If an EFS access point is specified in the authorizationConfig, the root directory parameter must either be omitted or set to /. This enforces the path that's set on the EFS access point.

#### transitEncryption

Type: String

Valid values: ENABLED | DISABLED

Required: No

Determines whether to enable encryption for Amazon EFS data that's in transit between the Amazon Batch host and the Amazon EFS server. Transit encryption must be enabled

if Amazon EFS IAM authorization is used. If this parameter is omitted, the default value of DISABLED is used. For more information, see Encrypting data in transit in the Amazon Elastic File System User Guide.

transitEncryptionPort

Type: Integer

Required: No

The port to use when sending encrypted data between the Amazon Batch host and the Amazon EFS server. If you don't specify a transit encryption port, it uses the port selection strategy that the Amazon EFS mount helper uses. The value must be between 0 and 65,535. For more information, see EFS Mount Helper in the Amazon Elastic File System User Guide.

authorizationConfig

Type: Object

Required: No

The authorization configuration details for the Amazon EFS file system.

accessPointId

Type: String

Required: No

The access point ID to use. If an access point is specified, the root directory value in the efsVolumeConfiguration must either be omitted or set to /. This enforces the path that's set on the EFS access point. If an access point is used, transit encryption must be enabled in the EFSVolumeConfiguration. For more information, see Working with Amazon EFS Access Points in the Amazon Elastic File System User Guide.

iam

Type: String

Valid values: ENABLED | DISABLED

Required: No

Determines whether to use the Amazon Batch job IAM role that's defined in a job definition when mounting the Amazon EFS file system. If enabled, transit encryption

must be enabled in the EFSVolumeConfiguration. If this parameter is omitted, the default value of DISABLED is used. For more information about IAM execution role, see Amazon Batch IAM execution role.

# Reference: Job definition examples

The job definition examples in the following topics illustrate how to use common patterns such as environment variables, parameter substitution, and volume mounts.

### **Contents**

- Reference: Environment variables
- Reference: Parameter substitution
- Reference: Test GPU functionality
- Reference: Multi-node parallel job

### Reference: Environment variables

The following example job definition uses environment variables to specify a file type and Amazon S3 URL. This particular example is from the <u>Creating a Simple "Fetch & Run" Amazon Batch Job</u> compute blog post. The <u>fetch\_and\_run.sh</u> script that's described in the blog post uses these environment variables to download the myjob.sh script from S3 and declare its file type.

Even though the command and environment variables are hardcoded into the job definition in this example, you can specify command and environment variable overrides to make the job definition more versatile.

```
"value": "2"
            }
        ],
        "command": [
            "myjob.sh",
            "60"
        "jobRoleArn": "arn:aws-cn:iam::123456789012:role/AWSBatchS3ReadOnly",
        "environment": Γ
            {
                "name": "BATCH_FILE_S3_URL",
                "value": "s3://amzn-s3-demo-source-bucket/myjob.sh"
            },
            {
                 "name": "BATCH_FILE_TYPE",
                "value": "script"
            }
        ],
        "user": "nobody"
    }
}
```

### **Reference: Parameter substitution**

The following example job definition illustrates how to allow for parameter substitution and to set default values.

The Ref:: declarations in the command section are used to set placeholders for parameter substitution. When you submit a job with this job definition, you specify the parameter overrides to fill in those values, such as the inputfile and outputfile. The parameters section that follows sets a default for codec, but you can override that parameter as needed.

For more information, see Parameters.

Reference: Parameter substitution 279

```
"type": "MEMORY",
                 "value": "2000"
            },
             {
                 "type": "VCPU",
                 "value": "2"
            }
        ],
        "command": [
            "ffmpeg",
            "-i",
             "Ref::inputfile",
            "-c",
             "Ref::codec",
             "-o",
             "Ref::outputfile"
        ],
        "jobRoleArn": "arn:aws-cn:iam::123456789012:role/ECSTask-S3FullAccess",
        "user": "nobody"
    }
}
```

# **Reference: Test GPU functionality**

The following example job definition tests if the GPU workload AMI described in <u>Use a GPU</u> workload AMI is configured properly. This example job definition runs the TensorFlow deep MNIST classifier example from GitHub.

You can create a file with the preceding JSON text called tensorflow\_mnist\_deep.json and then register an Amazon Batch job definition with the following command:

```
aws batch register-job-definition --cli-input-json file://tensorflow_mnist_deep.json
```

# Reference: Multi-node parallel job

The following example job definition illustrates a multi-node parallel job. For more information, see <u>Building a tightly coupled molecular dynamics workflow with multi-node parallel jobs in Amazon Batch in the Amazon Compute blog.</u>

```
"jobDefinitionName": "gromacs-jobdef",
  "jobDefinitionArn": "arn:aws-cn:batch:us-east-2:123456789012:job-definition/gromacs-
jobdef:1",
  "revision": 6,
  "status": "ACTIVE",
  "type": "multinode",
  "parameters": {},
  "nodeProperties": {
    "numNodes": 2,
    "mainNode": 0,
    "nodeRangeProperties": [
      {
        "targetNodes": "0:1",
        "container": {
          "image": "123456789012.dkr.ecr.us-east-2.amazonaws.com/gromacs_mpi:latest",
          "resourceRequirements": [
              {
                  "type": "MEMORY",
                  "value": "24000"
              },
                  "type": "VCPU",
```

# **Jobs**

Jobs are the unit of work that's started by Amazon Batch. Jobs can be invoked as containerized applications that run on Amazon ECS container instances in an ECS cluster.

Containerized jobs can reference a container image, command, and parameters. For more information, see JobDefinition.

You can submit a large number of independent, simple jobs.

### **Topics**

- Tutorial: submit a job
- Job states
- Amazon Batch job environment variables
- Automated job retries
- Job dependencies
- Job timeouts
- Amazon EKS jobs
- Multi-node parallel jobs
- Multi-node parallel jobs on Amazon EKS
- Array jobs
- Run GPU jobs
- Search Amazon Batch jobs in a job queue
- Networking modes for Amazon Batch jobs
- Tutorial: Send Amazon Batch job logs to CloudWatch Logs
- Tutorial: Review Amazon Batch job information

# Tutorial: submit a job

After you register a job definition, you can submit it as a job to an Amazon Batch job queue. You can override many of the parameters that are specified in the job definition at runtime.

Tutorial: submit a job 283

### To submit a job

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Jobs**.
- 4. Choose **Submit new job**.
- 5. For **Name**, enter a unique name for your job definition. The name can be up to 128 characters in length. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).
- 6. For **Job definition**, choose an existing job definition for your job. For more information, see Create a single-node job definition .
- 7. For **Job queue**, choose an existing job queue. For more information, see Create a job queue.
- 8. For **Job dependencies**, choose **Add Job dependencies**.
  - For **Job id**, enter the job ID for any dependencies. Then choose **Add job dependencies**. A job can have up to 20 dependencies. For more information, see **Job dependencies**.
- 9. (Array jobs only) For **Array size**, specify an array size between 2 and 10,000.
- 10. (Optional) Expand **Tags** and then choose **Add tag** to add tags to the resource. Enter a key and optional value, then choose **Add tag**.
- 11. Choose **Next page**.
- 12. In the **Job overrides** section:
  - a. (Optional) For **Scheduling priority**, enter a scheduling priority value between 0 and 100.
     Higher values are given higher priority.
  - b. (Optional) For **Job attempts**, enter the maximum number of times that Amazon Batch attempts to move the job to a RUNNABLE status. You can enter a number between 1 and 10. For more information, see <u>Automated job retries</u>.
  - c. (Optional) For Execution timeout, enter the timeout value (in seconds). The execution timeout is the length of time before an unfinished job is terminated. If an attempt exceeds the timeout duration, it's stopped and moves to a FAILED status. For more information, see Job timeouts. The minimum value is 60 seconds.

Tutorial: submit a job 284

### Important

Don't rely on jobs that run on Fargate resources to run for more than 14 days. After 14 days, the Fargate resources might no longer be available with the job being likely terminated.

- d. (Optional) Turn on **Propagate tags** to propagate tags from the job and job definition to the Amazon ECS task.
- 13. Expand Additional configuration.
- 14. (Optional) For Retry strategy conditions, choose Add evaluate on exit. Enter at least one parameter value and then choose an Action. For each set of conditions, Action must be set to either **Retry** or **Exit**. These actions mean the following:
  - Retry Amazon Batch retries until the number of job attempts that you specified is reached.
  - Exit Amazon Batch stops retrying the job.

### 

If you choose Add evaluate on exit, configure at least one parameter and choose either an Action or choose Remove evaluate on exit.

- 15. For **Parameters**, choose **Add parameters** to add parameter substitution placeholders. Then, enter a **Key** and an optional **Value**.
- 16. In the **Container overrides** section:
  - For **Command**, enter the commands into the field as their **JSON** string array equivalent.

This parameter maps to Cmd in the Create a container section of the Docker Remote API and the COMMAND parameter to docker run. For more information about the Docker CMD parameter, see https://docs.docker.com/engine/reference/builder/#cmd.



### Note

This parameter can't contain an empty string.

Tutorial: submit a job 285

For **vCPUs**, enter the number of vCPUs to reserve for the container. This parameter maps b. to CpuShares in the Create a container section of the Docker Remote API and the -cpu-shares option to docker run. Each vCPU is equivalent to 1,024 CPU shares. You must specify at least one vCPU.

For **Memory**, enter the memory limit that's available to the container. If your container attempts to exceed the memory specified here, the container is stopped. This parameter maps to Memory in the Create a container section of the Docker Remote API and the -memory option to docker run. You must specify at least 4 MiB of memory for a job.



### Note

To maximize your resource utilization, prioritize memory for jobs of a specific instance type. For more information, see Compute resource memory management.

- (Optional) For **Number of GPUs**, choose the number of GPUs to reserve for the container. d.
- (Optional) For Environment variables, choose Add environment variable to add e. environment variables as name-value pairs. These variables are passed to the container.
- f. Choose **Next page**.
- For **Job review**, review the configuration steps. If you need to make changes, choose **Edit**. g. When you're finished, choose **Create job definition**.

### Job states

When you submit a job to an Amazon Batch job queue, the job enters the SUBMITTED state. It then passes through the following states until it succeeds (exits with code 0) or fails (exits with a nonzero code). Amazon Batch jobs can have the following states:

#### **SUBMITTED**

A job that's submitted to the queue, and has not yet been evaluated by the scheduler. The scheduler evaluates the job to determine if it has any outstanding dependencies on the successful completion of any other jobs. If there are dependencies, the job is moved to PENDING. If there are no dependencies, the job is moved to RUNNABLE.

#### **PENDING**

A job that resides in the gueue and isn't yet able to run due to a dependency on another job or resource. After the dependencies are satisfied, the job is moved to RUNNABLE.

Job states 286

#### RUNNABLE

A job that resides in the queue, has no outstanding dependencies, and is therefore ready to be scheduled to a host. Jobs in this state are started as soon as sufficient resources are available in one of the compute environments that are mapped to the job's queue. However, jobs can remain in this state indefinitely when sufficient resources are unavailable.



#### Note

If your jobs don't progress to STARTING, see Jobs stuck in a RUNNABLE status in the troubleshooting section.

#### STARTING

These jobs have been scheduled to a host and the relevant container initiation operations are underway. After the container image is pulled and the container is up and running, the job transitions to RUNNING.

Image pull duration, Amazon EKS initContainer completion duration, and Amazon ECS containerDependency resolution duration occur in the STARTING state. The amount of time it takes to pull an image for your job is equivalent to the amount of time your job will be in the STARTING state.

For example, if it takes three minutes to pull the image for your job, your job will be in the STARTING state for three minutes. If initContainers takes a total of ten minutes to complete, then your Amazon EKS job will be in STARTING for ten minutes. If you have Amazon ECS containerDependencies sets in your Amazon ECS job, the job will be in STARTING until all container dependencies (their runtime) are resolved. STARTING is not included in timeouts; duration starts at RUNNING. For more information, see Job states.

### RUNNING

The job is running as a container job on an Amazon ECS container instance within a compute environment. When the job's container exits, the process exit code determines whether the job succeeded or failed. An exit code of 0 indicates success, and any non-zero exit code indicates failure. If the job associated with a failed attempt has any remaining attempts left in its optional retry strategy configuration, the job is moved to RUNNABLE again. For more information, see Automated job retries.

Job states 287



### Note

Logs for RUNNING jobs are available in CloudWatch Logs. The log group is /aws/batch/job, and the log stream name format is as follows: first200CharsOfJobDefinitionName/default/ecs\_task\_id. This format might change in the future.

After a job reaches the RUNNING status, you can programmatically retrieve its log stream name with the DescribeJobs API operation. For more information, see View Log Data Sent to CloudWatch Logs in the Amazon CloudWatch Logs User Guide. By default, these logs never expire. However, you can modify the retention period. For more information, see Change Log Data Retention in CloudWatch Logs in the Amazon CloudWatch Logs User Guide.

#### **SUCCEEDED**

The job has successfully completed with an exit code of 0. The job state for SUCCEEDED jobs is persisted in Amazon Batch for at least 7 days.



### Note

Logs for SUCCEEDED jobs are available in CloudWatch Logs. The log group is /aws/batch/job, and the log stream name format is as follows: first200CharsOfJobDefinitionName/default/ecs\_task\_id. This format may change in the future.

After a job reaches the RUNNING status, you can programmatically retrieve its log stream name with the DescribeJobs API operation. For more information, see View Log Data Sent to CloudWatch Logs in the Amazon CloudWatch Logs User Guide. By default, these logs never expires. However, you can modify the retention period. For more information, see Change Log Data Retention in CloudWatch Logs in the Amazon CloudWatch Logs User Guide.

#### **FAILED**

The job has failed all available attempts. The job state for FAILED jobs is persisted in Amazon Batch for at least 7 days.

Job states 288



### Note

Logs for FAILED jobs are available in CloudWatch Logs. The log group is /aws/batch/job, and the log stream name format is as follows: first200CharsOfJobDefinitionName/default/ecs\_task\_id. This format may change in the future.

After a job reaches the RUNNING status, you can programmatically retrieve its log stream with the DescribeJobs API operation. For more information, see View Log Data Sent to CloudWatch Logs in the Amazon CloudWatch Logs User Guide. By default, these logs never expire. However, you can modify the retention period. For more information, see Change Log Data Retention in CloudWatch Logs in the Amazon CloudWatch Logs User Guide.

# **Amazon Batch job environment variables**

Amazon Batch sets specific environment variables in container jobs. These environment variables provide introspection for the containers inside jobs. You can use the values of these variables in the logic of your applications. All variables that Amazon Batch set start with the AWS\_BATCH\_ prefix. This is a protected environment variable prefix. You can't use this prefix for your own variables in job definitions or overrides.

The following environment variables are available in job containers:

AWS\_BATCH\_CE\_NAME

This variable is set to the name of the compute environment where your job is placed.

AWS\_BATCH\_JOB\_ARRAY\_INDEX

This variable is only set in child array jobs. The array job index begins at 0, and each child job receives a unique index number. For example, an array job with 10 children has index values of 0-9. You can use this index value to control how your array job children are differentiated. For more information, see Tutorial: Use the array job index to control job differentiation.

AWS\_BATCH\_JOB\_ARRAY\_SIZE

This variable is set to the size of the parent array job. The size of the parent array job is passed to the child array job in this variable.

Job environment variables 289

### AWS\_BATCH\_JOB\_ATTEMPT

This variable is set to the job attempt number. The first attempt is numbered 1. For more information, see Automated job retries.

This variable is set to the Amazon Batch job ID.

This variable is set as the Kubernetes UID of the node object that's in the Kubernetes cluster that the pod runs on. This variable is only set for jobs that run on Amazon EKS resources. For more information, see <u>UIDs</u> in the *Kubernetes documentation*.

This variable is only set in multi-node parallel jobs. This variable is set to the index number of the job's main node. Your application code can compare the AWS\_BATCH\_JOB\_MAIN\_NODE\_INDEX to the AWS\_BATCH\_JOB\_NODE\_INDEX on an individual node to determine if it's the main node.

This variable is only set in multi-node parallel job child nodes. This variable isn't present on the main node, but is set to the private IPv4 address of the job's main node. Your child node's application code can use this address to communicate with the main node.

This variable is only set in multi-node parallel jobs. This variable is set to the node index number of the node. The node index begins at 0, and each node receives a unique index number. For example, a multi-node parallel job with 10 children has index values of 0-9.

This variable is only set in multi-node parallel jobs. This variable is set to the number of nodes that you requested for your multi-node parallel job.

This variable is set to the name of the job queue that your job was submitted to.

Job environment variables 290

# **Automated job retries**

You can apply a retry strategy to your jobs and job definitions that allows failed jobs to be automatically retried. Possible failure scenarios include the following:

- Any non-zero exit code from a container job
- Amazon EC2 instance failure or termination
- Internal Amazon service error or outage

When a job is submitted to a job queue and placed into the RUNNING state that's considered an attempt. By default, each job is given one attempt to move to either the SUCCEEDED or FAILED job state. However, both the job definition and the job submission workflows can be used to specify a retry strategy with between 1 and 10 attempts. If <a href="evaluateOnExit">evaluateOnExit</a> is specified, it can contain up to 5 retry strategies. If <a href="evaluateOnExit">evaluateOnExit</a> is specified, but none of the retry strategies match, then the job is retried. For jobs that don't match to exit, add a final entry that exits for any reason. For example, this <a href="evaluateOnExit">evaluateOnExit</a> object has two entries that with actions of RETRY and a final entry with an action of EXIT.

At runtime, the AWS\_BATCH\_JOB\_ATTEMPT environment variable is set to the container's corresponding job attempt number. The first attempt is numbered 1, and subsequent attempts are in ascending order (for example, 2, 3, 4).

Automated job retries 291

For example, suppose that a job attempt fails for any reason and the number of attempts specified in the retry configuration is greater than the AWS BATCH JOB ATTEMPT number. Then, the job is placed back in the RUNNABLE state. For more information, see Job states.



#### Note

Jobs that are cancelled or terminated aren't retried. Also, jobs that fail because of an invalid job definition aren't retried.

For more information, see Retry strategy, Create a single-node job definition, Tutorial: submit a job and Stopped tasks error codes.

# Job dependencies

When you submit an Amazon Batch job, you can specify the job IDs that the job depends on. When you do so, the Amazon Batch scheduler ensures that your job is run only after the specified dependencies have successfully completed. After they succeed, the dependent job transitions from PENDING to RUNNABLE and then to STARTING and RUNNING. If any of the job dependencies fail, the dependent job automatically transitions from PENDING to FAILED.

For example, Job A can express a dependency on up to 20 other jobs that must succeed before it can run. You can then submit additional jobs that have a dependency on Job A and up to 19 other jobs.

For array jobs, you can specify a SEQUENTIAL type dependency without specifying a job ID so that each child array job completes sequentially, starting at index 0. You can also specify an N\_TO\_N type dependency with a job ID. That way, each index child of this job must wait for the corresponding index child of each dependency to complete before it can begin. For more information, see Array jobs.

To submit an Amazon Batch job with dependencies, see Tutorial: submit a job.

Resource-aware scheduling lets you schedule jobs based on consumable resources that are needed to run your jobs. You specify the consumable resources which are needed for a job to run, and Batch takes these resource dependencies into account when it schedules a job. You can reduce the underutilization of compute resources by allocating only the jobs that have all the required resources available. Resource-aware scheduling is available for both FIFO and fair-share scheduling

Job dependencies 292

policies and can be used with all compute platforms supported by Batch including EKS, ECS, and Fargate. It can be used with Array jobs, Multi-node parallel (MNP) jobs, and with regular Batch jobs.

# **Job timeouts**

You can configure a timeout duration for your jobs so that if a job runs longer than that, Amazon Batch terminates the job. For example, you might have a job that you know should only take 15 minutes to complete. Sometimes your application gets stuck in a loop and runs forever, so you can set a timeout of 30 minutes to terminate the stuck job.

### Important

By default, Amazon Batch doesn't have a job timeout. If you don't define a job timeout, the job runs until the container exits.

You specify an attemptDurationSeconds parameter, which must be at least 60 seconds, either in your job definition, or when you submit the job. When this number of seconds has passed following the job attempt's startedAt timestamp, Amazon Batch terminates the job. On the compute resource, your job's container receives a SIGTERM signal to give your application a chance to shut down gracefully. If the container is still running after 30 seconds, a SIGKILL signal is sent to forcefully shut down the container.

Timeout terminations are handled on a best-effort basis. You shouldn't expect your timeout termination to happen exactly when the job attempt times out (it may take a few seconds longer). If your application requires precise timeout execution, you should implement this logic within the application. If you have a large number of jobs timing out concurrently, the timeout terminations behave as a first in, first out queue, where jobs are terminated in batches.



### Note

There's no maximum timeout value for an Amazon Batch job.

If a job is terminated for exceeding the timeout duration, it isn't retried. If a job attempt fails on its own, then it can retry if retries are enabled, and the timeout countdown is started over for the new attempt.

Job timeouts 293

### 

Jobs that run on Fargate resources can't expect to run for more than 14 days. If the timeout duration exceeds 14 days, the Fargate resources may no longer be available and the job will be terminated.

For array jobs, child jobs have the same timeout configuration as the parent job.

For information about submitting an Amazon Batch job with a timeout configuration, see Tutorial: submit a job.

# Amazon EKS jobs

A job is the smallest unit of work in Amazon Batch. An Amazon Batch job on Amazon EKS has a one-to-one mapping to a Kubernetes pod. An Amazon Batch job definition is a template for an Amazon Batch job. When you submit an Amazon Batch job, you reference a job definition, target a job queue, and provide a name for a job. In the job definition of an Amazon Batch job on Amazon EKS, the eksProperties parameter defines the set of parameters that an Amazon Batch on Amazon EKS job supports. In a SubmitJob request, the eksPropertiesOverride parameter allows for overrides to some common parameters. This way, you can use templates of job definitions for multiple jobs. When a job is dispatched to your Amazon EKS cluster, Amazon Batch transforms the job into a podspec (Kind: Pod). The podspec uses some additional Amazon Batch parameters to ensure that jobs are scaled and scheduled correctly. Amazon Batch combines labels and taints to ensure jobs run only on Amazon Batch managed nodes and that other pods don't run on those nodes.

### Important

- If the hostNetwork parameter isn't explicitly set in an Amazon EKS job definition, the pod networking mode in Amazon Batch defaults to host mode. More specifically, the following settings are applied: hostNetwork=true and dnsPolicy=ClusterFirstWithHostNet.
- Amazon Batch cleans up job pods soon after a pod completes its job. To see pod application logs, configure a logging service for your cluster. For more information, see Use CloudWatch Logs to monitor Amazon Batch on Amazon EKS jobs.

Amazon EKS jobs 294

### **Topics**

- Tutorial: Map a running job to a pod and a node
- Tutorial: Map a running pod back to its job

# Tutorial: Map a running job to a pod and a node

The podProperties of a running job have podName and nodeName parameters set for the current job attempt. Use the DescribeJobs API operation to view these parameters.

The following is example output.

```
$ aws batch describe-jobs --job 2d044787-c663-4ce6-a6fe-f2baf7e51b04
{
 "jobs": [
   "status": "RUNNING",
   "jobArn": "arn:aws-cn:batch:us-east-1:123456789012:job/2d044787-c663-4ce6-a6fe-
f2baf7e51b04",
   "jobDefinition": "arn:aws-cn:batch:us-east-1:123456789012:job-definition/
MyJobOnEks_SleepWithRequestsOnly:1",
   "jobQueue": "arn:aws-cn:batch:us-east-1:123456789012:job-queue/My-Eks-JQ1",
   "jobId": "2d044787-c663-4ce6-a6fe-f2baf7e51b04",
   "eksProperties": {
    "podProperties": {
     "nodeName": "ip-192-168-55-175.ec2.internal",
     "containers": [
      {
       "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
       "resources": {
        "requests": {
         "cpu": "1",
         "memory": "1024Mi"
        }
       }
      }
     ],
     "podName": "aws-batch.b0aca953-ba8f-3791-83e2-ed13af39428c"
   }
  }
 ]
```

}

For a job with retries enabled, the podName and nodeName of every completed attempt is in the eksAttempts list parameter of the <u>DescribeJobs</u> API operation. The podName and nodeName of the current running attempt are in the podProperties object.

# Tutorial: Map a running pod back to its job

A pod has labels that indicate the jobId and uuid of the compute environment that it belongs to. Amazon Batch injects environment variables so the job's runtime can reference job information. For more information, see <a href="Manazon Batch job environment variables"><u>Amazon Batch job environment variables</u></a>. You can view this information by running the following command. The output is as follows.

```
$ kubectl describe pod aws-batch.14638eb9-d218-372d-ba5c-1c9ab9c7f2a1 -n my-aws-batch-
namespace
              aws-batch.14638eb9-d218-372d-ba5c-1c9ab9c7f2a1
Name:
Namespace:
              my-aws-batch-namespace
Priority:
Node:
              ip-192-168-45-88.ec2.internal/192.168.45.88
Start Time:
              Wed, 26 Oct 2022 00:30:48 +0000
Labels:
              batch.amazonaws.com/compute-environment-uuid=5c19160b-
d450-31c9-8454-86cf5b30548f
              batch.amazonaws.com/job-id=f980f2cf-6309-4c77-a2b2-d83fbba0e9f0
              batch.amazonaws.com/node-uid=a4be5c1d-9881-4524-b967-587789094647
Status:
              Running
IP:
              192,168,45,88
IPs:
  IP: 192.168.45.88
Containers:
  default:
                   public.ecr.aws/amazonlinux/amazonlinux:2
    Image:
    Environment:
      AWS_BATCH_JOB_KUBERNETES_NODE_UID:
                                           a4be5c1d-9881-4524-b967-587789094647
                                           f980f2cf-6309-4c77-a2b2-d83fbba0e9f0
      AWS_BATCH_JOB_ID:
      AWS_BATCH_JQ_NAME:
                                           My-Eks-JQ1
      AWS_BATCH_JOB_ATTEMPT:
                                           1
      AWS_BATCH_CE_NAME:
                                           My-Eks-CE1
. . .
```

### Features that Amazon Batch Amazon EKS jobs support

These are the Amazon Batch specific features that are also common to Kubernetes jobs that run on Amazon EKS:

- Job dependencies
- Array jobs
- Job timeouts
- Automated job retries
- Use fair-share scheduling to help schedule jobs

#### KubernetesSecrets and ServiceAccounts

Amazon Batch supports referencing Kubernetes Secrets and ServiceAccounts. You can configure pods to use Amazon EKS IAM roles for service accounts. For more information, see Configuring pods to use a Kubernetes service account in the *Amazon EKS User Guide*.

#### **Related documents**

- Memory and vCPU considerations for Amazon Batch on Amazon EKS
- Run GPU jobs
- Jobs stuck in a RUNNABLE status

# Multi-node parallel jobs

You can use multi-node parallel jobs to run single jobs that span multiple Amazon EC2 instances. With Amazon Batch multi-node parallel jobs (also known as *gang scheduling*), you can run large-scale, high-performance computing applications and distributed GPU model training without the need to launch, configure, and manage Amazon EC2 resources directly. An Amazon Batch multi-node parallel job is compatible with any framework that supports IP-based, inter-node communication. Examples include Apache MXNet, TensorFlow, Caffe2, or Message Passing Interface (MPI).

Multi-node parallel jobs are submitted as a single job. However, your job definition (or job submission node overrides) specifies the number of nodes to create for the job and what node groups to create. Each multi-node parallel job contains a **main node**, which is launched first. After

Multi-node parallel jobs 297

the main node is up, the child nodes are launched and started. The job is finished only if the main node exits. All child nodes are then stopped. For more information, see Node groups.

Multi-node parallel job nodes are single-tenant. This means that only a single job container is run on each Amazon EC2 instance.

The final job status (SUCCEEDED or FAILED) is determined by the final job status of the main node. To get the status of a multi-node parallel job, describe the job by using the job ID that was returned when you submitted the job. If you need the details for child nodes, describe each child node individually. You can address nodes using the #N notation (starting with 0). For example, to access the details of the second node of a job, describe <a href="https://www.batch\_job\_id#1">aws\_batch\_job\_id#1</a> using the Amazon Batch <a href="DescribeJobs">DescribeJobs</a> API operation. The started, stoppedAt, statusReason, and exit information for a multi-node parallel job is populated from the main node.

If you specify job retries, a main node failure causes another attempt to occur. Child node failures don't cause more attempts to occur. Each new attempt of a multi-node parallel job updates the corresponding attempt of its associated child nodes.

To run multi-node parallel jobs on Amazon Batch, your application code must contain the frameworks and libraries that are necessary for distributed communication.

### **Topics**

- Environment variables
- Node groups
- Job lifecycle for MNP jobs
- Compute environment considerations for MNP with Amazon Batch

### **Environment variables**

At runtime, each node is configured the standard environment variables that all Amazon Batch jobs receive. In addition, the nodes are configured with the following environment variables that are specific to multi-node parallel jobs:

AWS\_BATCH\_JOB\_MAIN\_NODE\_INDEX

This variable is set to the index number of the job's main node. Your application code can compare the AWS\_BATCH\_JOB\_MAIN\_NODE\_INDEX to the AWS\_BATCH\_JOB\_NODE\_INDEX on an individual node to determine if it's the main node.

Environment variables 298

### AWS BATCH JOB MAIN NODE PRIVATE IPV4 ADDRESS

This variable is only set in multi-node parallel job child nodes. This variable isn't present on the main node. This variable is set to the private IPv4 address of the job's main node. Your child node's application code can use this address to communicate with the main node.

This variable is set to the node index number of the node. The node index begins at 0, and each node receives a unique index number. For example, a multi-node parallel job with 10 children has index values of 0-9.

This variable is set to the number of nodes that you have requested for your multi-node parallel job.

# **Node groups**

A node group is an identical group of job nodes that all share the same container properties. You can use Amazon Batch to specify up to five distinct node groups for each job.

Each group can have its own container images, commands, environment variables, and so on. For example, you can submit a job that requires a single c5.xlarge instance for the main node and five c5.xlarge instance child nodes. Each of these distinct node groups may specify different container images or commands to run for each job.

Alternatively, all of the nodes in your job can use a single node group. Moreover, your application code can differentiate node roles such as the main node and child node. It does this by comparing the AWS BATCH JOB MAIN NODE INDEX environment variable against its own value for AWS\_BATCH\_JOB\_NODE\_INDEX. You can have up to 1,000 nodes in a single job. This is the default limit for instances in an Amazon ECS cluster. You can request to increase this limit.



### Note

Currently all node groups in a multi-node parallel job must use the same instance type.

Node groups 299

# Job lifecycle for MNP jobs

When you submit a multi-node parallel job, the job enters the SUBMITTED status. Then, the job waits for any job dependencies to finish. The job also moves to the RUNNABLE status. Last, Amazon Batch provisions the instance capacity that's required to run your job and launches these instances.

Each multi-node parallel job contains a **main node**. The main node is a single subtask that Amazon Batch monitors to determine the outcome of the submitted multi node job. The main node is launched first and it moves to the STARTING status. The timeout value specified in the attemptDurationSeconds parameter applies to the whole job and not to the nodes.

When the main node reaches the RUNNING status after the node's container is running, the child nodes are launched and they also move to the STARTING status. The child nodes come up in random order. There are no guarantees on the timing or ordering of child node launch. To ensure that the all the nodes of the jobs are in the RUNNING status after the node's container is running, your application code can query the Amazon Batch API to get the main node and child node information. Alternatively, the application code can wait until all nodes are online before starting any distributed processing task. The private IP address of the main node is available as the AWS\_BATCH\_JOB\_MAIN\_NODE\_PRIVATE\_IPV4\_ADDRESS environment variable in each child node. Your application code may use this information to coordinate and communicate data between each task.

As individual nodes exit, they move to SUCCEEDED or FAILED, depending on their exit code. If the main node exits, the job is considered finished, and all of the child nodes are stopped. If a child node dies, Amazon Batch doesn't take any action on the other nodes in the job. If you don't want your job to continue with a reduced number of nodes, you must factor this into your application code. Doing this terminates or cancels the job.

# **Compute environment considerations for MNP with Amazon Batch**

There are several things to consider when configuring compute environments to run multi-node parallel jobs with Amazon Batch.

- Multi-node parallel jobs aren't supported on UNMANAGED compute environments.
- If you want to submit multi-node parallel jobs to a compute environment, create a *cluster* placement group in a single Availability Zone and associate it with your compute resources. This keeps your multi-node parallel jobs on a logical grouping of instances close with high network flow potential. For more information, see Placement Groups in the *Amazon EC2 User Guide*.

Job lifecycle 300

• Multi-node parallel jobs aren't supported on compute environments that use Spot Instances.

- Amazon Batch multi-node parallel jobs use the Amazon ECS awsvpc network mode, which
  gives your multi-node parallel job containers the same networking properties as Amazon EC2
  instances. Each multi-node parallel job container gets its own elastic network interface, a
  primary private IP address, and an internal DNS hostname. The network interface is created in
  the same VPC subnet as its host compute resource.
- Your compute environment might have no more than five security groups associated with it. The
  elastic network interfaces that are created and attached to an MNP task will use the security
  groups specified in your compute environment, If you don't specify a security group, the default
  security group for the VPC is used.
- The aws vpc network mode doesn't provide the elastic network interfaces for multi-node parallel jobs with public IP addresses. To access the internet, your compute resources must be launched in a private subnet that is configured to use a NAT gateway. For more information, see <a href="NAT">NAT</a>
   Gateways in the Amazon VPC User Guide. Inter-node communication must use the private IP address or DNS hostname for the node. Multi-node parallel jobs that run on compute resources within public subnets don't have outbound network access. To create a VPC with private subnets and a NAT gateway, see Create a virtual private cloud.
- The elastic network interfaces that are created and attached to your compute resources can't be detached manually or modified by your account. This is to prevent the accidental deletion of an elastic network interface that's associated with a running job. To release the elastic network interfaces for a task, terminate the job.
- Your compute environment must have enough maximum vCPUs to support your multi-node parallel job.
- Your Amazon EC2 instance quota include the number of instances that's required to run your job. For example, suppose that your job requires 30 instances, but your account can only run 20 instances in a Region. Then, your job will get stuck in RUNNABLE status.
- If you specify an instance type for a node group in a multi-node parallel job, your compute environment must launch that instance type.

# Multi-node parallel jobs on Amazon EKS

You can use Amazon Batch on Amazon Elastic Kubernetes Service to run multi-node parallel (MNP) jobs (also known as *gang scheduling*) on your managed Kubernetes clusters. This option is commonly used for large, tightly-coupled, high-performance jobs that can't be run on a single Amazon Elastic Compute Cloud instance. For more information, see Multi-node parallel jobs.

You can use this feature to run Amazon EKS managed Kubernetes-specific high-performance computing applications, large language model training, and other Artificial Intelligence (AI)/ Machine Learning (ML) jobs.

### **Topics**

- Running MNP jobs
- Create an Amazon EKS MNP job definition
- Submit an Amazon EKS MNP job
- Override an Amazon EKS MNP job definition

# **Running MNP jobs**

Amazon Batch supports MNP jobs on Amazon Elastic Container Service and Amazon EKS using Amazon EC2. The following provides more specifics about the instance and container parameters for the feature.

### **Instance quotas for MNP on Amazon EKS**

- Up to 1000 instances can be used for a single MNP job.
- Up to 5000 instances can join a single Amazon EKS cluster.
- Up to 5 compute environments can be clustered and attached to a job-queue.

For example, you can scale up to 5 clustered compute environments in a job queue and 1000 instances in each compute environment.

In addition to the instance parameters, it's important to note that you can't use Fargate for MNP jobs through either service.

You can use only one instance type in each MNP job. You can change the instance type by updating the compute environment, or when you define a new compute environment. You can also specify the instance type, and provide vCPU and memory requirements when creating the job definition.

### **Container quotas for MNP on Amazon EKS**

- A multi-node parallel job supports one pod per node.
- Up to 10 containers (or 10 init containers. For more information see <u>Init Containers</u> in the *Kubernetes documentation*.) in each pod.

Running MNP jobs 302

- Up to 5 node ranges in each MNP job.
- Up 10 distinct container images in each node range.

For example, you can run up to a maximum of 10,000 containers in a single MNP job that contains 5 node ranges and a total of 50 unique images.

### Running MNP jobs in a private Amazon VPC and an Amazon EKS cluster

MNP jobs can run on any Amazon EKS cluster whether it has public Internet or not. When using an Amazon EKS cluster with only private network access be sure that Amazon Batch can access the Amazon EKS control plane and the managed Kubernetes API server. You can grant the necessary access through Amazon Virtual Private Cloud endpoints. For more information, see <a href="Configure an endpoint service">Configure an endpoint service</a>.

Amazon EKS cluster Pods can't download an image from a public source since the private VPC doesn't have Internet access. Your Amazon EKS cluster must pull images from a container registry that's within your Amazon VPC. You can create an <a href="Mazon Elastic Container Registry"><u>Amazon ECR</u></a>) in your Amazon VPC and copy container images to it for your nodes access.

You can also create a pull through cache rule with Amazon ECR. Once a pull through cache rule is created for an external public registry, you can simply pull an image from that external public registry using your Amazon ECR private registry URI. Then Amazon ECR creates a repository and caches the image. When a cached image is pulled using the Amazon ECR private registry URI, Amazon ECR checks the remote registry to see if there is a new version of the image and will update your private registry up to one time every 24 hours. For more information, see <a href="Creating a pull through cache rule in Amazon ECR">Creating a pull through cache rule in Amazon ECR</a>.

### **Error notification**

If your MNP jobs are blocked, you can receive notifications through the Amazon Web Services Management Console and Amazon EventBridge. For example, if an MNP job is stuck at the head of the queue, you can be notified about the issue along with information about what caused it so that you can take prompt action to unblock your job queue. Optionally, you can auto-terminate the MNP job if no action is taken within a distinct amount of time, which can be defined in the job queue template. For more information, see Resource: Job queue blocked events

Running MNP jobs 303

# **Create an Amazon EKS MNP job definition**

To define and run MNP jobs on Amazon EKS, there are new parameters within the RegisterJobDefinition and SubmitJob API operations.

- Use eksProperties under the nodeProperties section to define your MNP job definition.
- Use <u>eksPropertiesOverride</u> under the <u>nodePropertyOverrides</u> section to override the parameters defined in the job definition when submitting an MNP job.

These actions can be defined through API operations and the Amazon Web Services Management Console.

### Reference: Register the Amazon EKS MNP job definition request payload

The following example illustrates how you can register an Amazon EKS MNP job definition with two nodes.

```
"jobDefinitionName": "MyEksMnpJobDefinition",
"type": "multinode",
"nodeProperties": {
  "numNodes": 2,
  "mainNode": 0,
  "nodeRangeProperties": [
      "targetNodes" : "0:",
      "eksProperties": {
        "podProperties": {
          "containers": [
              "name": "test-eks-container-1",
              "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
              "command": [
                "sleep",
                "60"
              ],
              "resources": {
                "limits": {
                  "cpu": "1",
                  "memory": "1024Mi"
                }
```

```
},
                 "securityContext":{
                   "runAsUser":1000,
                   "runAsGroup":3000,
                   "privileged":true,
                   "readOnlyRootFilesystem":true,
                   "runAsNonRoot":true
                }
              }
            ],
             "initContainers": [
                {
                   "name": "init-ekscontainer",
                   "image": "public.ecr.aws/amazonlinux/amazonlinux:2",
                   "command": [
                      "echo",
                      "helloWorld"
                    ],
                    "resources": {
                      "limits": {
                        "cpu": "1",
                        "memory": "1024Mi"
                      }
                   }
                }
            ],
             "metadata": {
                "labels": {
                   "environment" : "test"
                }
            }
          }
        }
      }
    ]
  }
}
```

To register the job definition using the Amazon CLI, copy the definition to a local file named *MyEksMnpJobDefinition.json* and run the following command.

```
aws batch register-job-definition --cli-input-json file://MyEksMnpJobDefinition.json
```

You will receive the following JSON response.

```
{
    "jobDefinitionName": "MyEksMnpJobDefinition",
    "jobDefinitionArn": "arn:aws:batch:us-east-1:0123456789:job-definition/
MyEksMnpJobDefinition:1",
    "revision": 1
}
```

# Submit an Amazon EKS MNP job

To submit a job using the registered job definition, enter the following command. Replace the value of <EKS\_JOB\_QUEUE\_NAME> with the name or ARN of a pre-existing job queue associated with an Amazon EKS compute environment.

```
aws batch submit-job --job-queue < EKS_JOB_QUEUE_NAME > \
    --job-definition MyEksMnpJobDefinition \
    --job-name myFirstEksMnpJob
```

You will receive the following JSON response.

```
{
    "jobArn": "arn:aws-cn:batch:region:account:job/9b979cce-9da0-446d-90e2-
ffa16d52af68",
    "jobName": "myFirstEksMnpJob",
    "jobId": "<JOB_ID>"
}
```

You can check the status of the job using the returned jobId with the following command.

```
aws batch describe-jobs --jobs <\sum_ID>
```

# Override an Amazon EKS MNP job definition

Optionally, you can override the job definition details (such as changing the MNP job size or child job details). The following provides an example JSON request payload to submit a five node MNP job, and changes to the test-eks-container-1 container's command.

```
{
```

```
"numNodes": 5,
  "nodePropertyOverrides": [
      "targetNodes": "0:",
      "eksPropertiesOverride": {
        "podProperties": {
           "containers": [
             {
               "name": "test-eks-container-1",
               "command": [
                 "sleep",
                 "150"
               ]
             }
          ]
        }
      }
    }
  ]
}
```

To submit a job with these overrides, save the example to a local file, *eks-mnp-job-nodeoverride.json*, and use the Amazon CLI to submit the job with the overrides.

# **Array jobs**

An array job is a job that shares common parameters, such as the job definition, vCPUs, and memory. It runs as a collection of related yet separate basic jobs that might be distributed across multiple hosts and might run concurrently. Array jobs are the most efficient way to run extremely parallel jobs such as Monte Carlo simulations, parametric sweeps, or large rendering jobs.

Amazon Batch array jobs are submitted just like regular jobs. However, you specify an array size (between 2 and 10,000) to define how many child jobs should run in the array. If you submit a job with an array size of 1000, a single job runs and spawns 1000 child jobs. The array job is a reference or pointer to manage all the child jobs. This way, you can submit large workloads with a single query. The timeout specified in the attemptDurationSeconds parameter applies to each child job. The parent array job does not have a timeout.

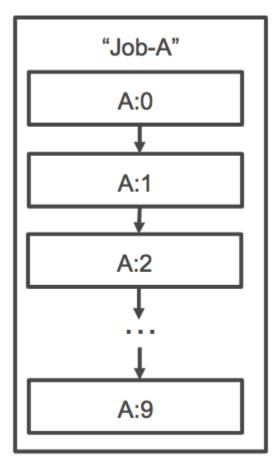
When you submit an array job, the parent array job gets a normal Amazon Batch job ID. Each child job has the same base ID. However, the array index for the child job is appended to the end of the parent ID, such as <code>example\_job\_ID:</code> of for the first child job of the array.

Array jobs 307

The parent array job can enter a SUBMITTED, PENDING, FAILED, or SUCCEEDED status. An array parent job is updated to PENDING when any child job is updated to RUNNABLE. For more information about job dependencies, see Job dependencies.

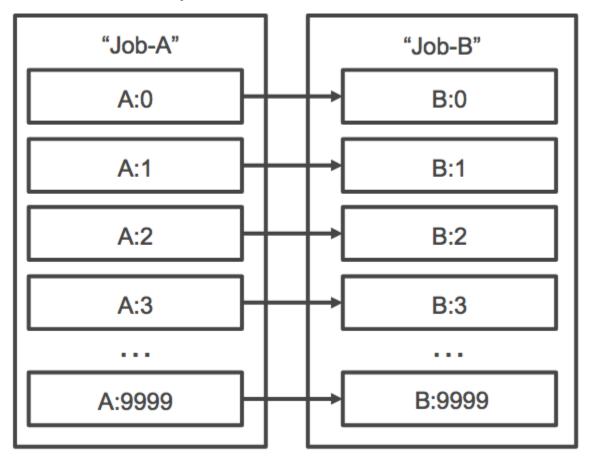
At runtime, the AWS\_BATCH\_JOB\_ARRAY\_INDEX environment variable is set to the container's corresponding job array index number. The first array job index is numbered 0, and subsequent attempts are in ascending order (for example, 1, 2, and 3). You can use this index value to control how your array job children are differentiated. For more information, see <a href="Internal: Use the array job"><u>Tutorial: Use the array job index to control job differentiation.</u></a>

For array job dependencies, you can specify a type for a dependency, such as SEQUENTIAL or N\_TO\_N. You can specify a SEQUENTIAL type dependency (without specifying a job ID) so that each child array job completes sequentially, starting at index 0. For example, if you submit an array job with an array size of 100, and specify a dependency with type SEQUENTIAL, 100 child jobs are spawned sequentially, where the first child job must succeed before the next child job starts. The figure below shows Job A, an array job with an array size of 10. Each job in Job A's child index is dependent on the previous child job. Job A:1 can't start until job A:0 finishes.



Array jobs 308

You can also specify an N\_TO\_N type dependency with a job ID for array jobs. That way, each index child of this job must wait for the corresponding index child of each dependency to complete before it can begin. The following figure shows Job A and Job B, two array jobs with an array size of 10,000 each. Each job in Job B's child index is dependent on the corresponding index in Job A. Job B:1 can't start until job A:1 finishes.



If you cancel or terminate a parent array job, all the child jobs are cancelled or terminated with it. You can cancel or terminate individual child jobs (which moves them to a FAILED status) without affecting the other child jobs. However, if a child array job fails (on its own, or by manually cancelling or terminating the job), the parent job also fails. In this scenario, the parent job transitions to FAILED when all child jobs complete.

### **Topics**

- Reference: Example of an array job workflow
- Tutorial: Use the array job index to control job differentiation

Array jobs 309

# Reference: Example of an array job workflow

A common workflow for Amazon Batch customers is to run a prerequisite setup job, run a series of commands against a large number of input tasks, and then conclude with a job that aggregates results and writes summary data to Amazon S3, DynamoDB, Amazon Redshift, or Aurora.

### For example:

• JobA: A standard, non-array job that performs a quick listing and metadata validation of objects in an Amazon S3 bucket, BucketA. The SubmitJob JSON syntax is as follows.

```
{
    "jobName": "JobA",
    "jobQueue": "ProdQueue",
    "jobDefinition": "JobA-list-and-validate:1"
}
```

 JobB: An array job with 10,000 copies that is dependent upon JobA that runs CPU-intensive commands against each object in BucketA and uploads results to BucketB. The <u>SubmitJob</u> JSON syntax is as follows.

```
}
    "jobName": "JobB",
    "jobQueue": "ProdQueue",
    "jobDefinition": "JobB-CPU-Intensive-Processing:1",
    "containerOverrides": {
        "resourceRequirements": [
                "type": "MEMORY",
                "value": "4096"
            },
                "type": "VCPU",
                "value": "32"
            }
        ]
    "arrayProperties": {
        "size": 10000
    },
    "depends0n": [
```

```
"jobId": "JobA_job_ID"
}
]
}
```

 JobC: Another 10,000 copy array job that's dependent upon JobB with an N\_T0\_N dependency model, that runs memory-intensive commands against each item in BucketB, writes metadata to DynamoDB, and uploads the resulting output to BucketC. The <u>SubmitJob</u> JSON syntax is as follows.

```
{
    "jobName": "JobC",
    "jobQueue": "ProdQueue",
    "jobDefinition": "JobC-Memory-Intensive-Processing:1",
    "containerOverrides": {
        "resourceRequirements": [
                 "type": "MEMORY",
                 "value": "32768"
            },
            {
                 "type": "VCPU",
                 "value": "1"
            }
        ]
    "arrayProperties": {
        "size": 10000
    },
    "depends0n": [
        {
             "jobId": "JobB_job_ID",
            "type": "N_TO_N"
        }
    ]
}
```

 JobD: An array job that performs 10 validation steps that each need to query DynamoDB and might interact with any of the above Amazon S3 buckets. Each of the steps in JobD run the same command. However, the behavior is different based on the value of the AWS\_BATCH\_JOB\_ARRAY\_INDEX environment variable within the job's container. These

validation steps run sequentially (for example, JobD: 0 and then JobD: 1). The <u>SubmitJob</u> JSON syntax is as follows.

```
{
    "jobName": "JobD",
    "jobQueue": "ProdQueue",
    "jobDefinition": "JobD-Sequential-Validation:1",
    "containerOverrides": {
        "resourceRequirements": [
                 "type": "MEMORY",
                 "value": "32768"
            },
            {
                 "type": "VCPU",
                 "value": "1"
            }
        ]
   }
    "arrayProperties": {
        "size": 10
    },
    "depends0n": [
        {
            "jobId": "JobC_job_ID"
        },
        {
            "type": "SEQUENTIAL"
        },
    ]
}
```

• JobE: A final, non-array job that performs some simple cleanup operations and sends an Amazon SNS notification with a message that the pipeline has completed and a link to the output URL. The SubmitJob JSON syntax is as follows.

```
{
    "jobName": "JobE",
    "jobQueue": "ProdQueue",
    "jobDefinition": "JobE-Cleanup-and-Notification:1",
    "parameters": {
        "SourceBucket": "s3://amzn-s3-demo-source-bucket",
```

# Tutorial: Use the array job index to control job differentiation

This tutorial describes how to use the AWS\_BATCH\_JOB\_ARRAY\_INDEX environment variable to differentiate the child jobs. Each child job is assigned to this variable. The example uses the child job's index number to read a specific line in a file. Then, it substitutes the parameter associated with that line number with a command inside the job's container. The result is that you can have multiple Amazon Batch jobs that run the same Docker image and command arguments. However, the results are different because the array job index is used as a modifier.

In this tutorial, you create a text file that has all of the colors of the rainbow, each on its own line. Then, you create an entrypoint script for a Docker container that converts the index into a value that can be used for a line number in the color file. The index starts at zero, but line numbers start at one. Create a Dockerfile that copies the color and index files to the container image and sets ENTRYPOINT for the image to the entrypoint script. The Dockerfile and resources are built to a Docker image that's pushed to Amazon ECR. You then register a job definition that uses your new container image, submit an Amazon Batch array job with that job definition, and view the results.

#### **Topics**

- Prerequisites
- Tutorial: Build a container image
- Tutorial: Push your image to Amazon ECR
- Tutorial: Create and register a job definition
- Tutorial: Submit an Amazon Batch array job
- Tutorial: View your array job logs

# **Prerequisites**

This tutorial workflow has the following prerequisites:

 An Amazon Batch compute environment. For more information, see Create a compute environment.

- An Amazon Batch job queue and associated compute environment. For more information, see Create a job queue.
- The Amazon CLI installed on your local system. For more information, see Installing the Amazon Command Line Interface in the Amazon Command Line Interface User Guide.
- Docker installed on your local system. For more information, see About Docker CE in the Docker documentation.

### **Tutorial: Build a container image**

You can use the AWS\_BATCH\_JOB\_ARRAY\_INDEX in a job definition in the command parameter. However, we recommend that you create a container image that uses the variable in an entrypoint script instead. This section describes how to create such a container image.

#### To build your Docker container image

- 1. Create a new directory to use as your Docker image workspace and navigate to it.
- Create a file named colors.txt in your workspace directory and paste the following into it. 2.

```
red
orange
yellow
green
blue
indigo
violet
```

Create a file named print-color. sh in your workspace directory and paste the following into it.



#### Note

The LINE variable is set to the AWS\_BATCH\_JOB\_ARRAY\_INDEX + 1 because the array index starts at 0, but line numbers start at 1. The COLOR variable is set to the color in colors.txt that's associated with its line number.

```
#!/bin/sh
LINE=$((AWS_BATCH_JOB_ARRAY_INDEX + 1))
COLOR=$(sed -n ${LINE}p /tmp/colors.txt)
echo My favorite color of the rainbow is $COLOR.
```

4. Create a file named Dockerfile in your workspace directory and paste the following content into it. This Dockerfile copies the previous files to your container and sets the entrypoint script to run when the container starts.

```
FROM busybox
COPY print-color.sh /tmp/print-color.sh
COPY colors.txt /tmp/colors.txt
RUN chmod +x /tmp/print-color.sh
ENTRYPOINT /tmp/print-color.sh
```

5. Build your Docker image.

```
$ docker build -t print-color .
```

6. Test your container with the following script. This script sets the AWS\_BATCH\_JOB\_ARRAY\_INDEX variable to 0 locally and then increments it to simulate what an array job with seven children does.

```
$ AWS_BATCH_JOB_ARRAY_INDEX=0
while [ $AWS_BATCH_JOB_ARRAY_INDEX -le 6 ]
do
    docker run -e AWS_BATCH_JOB_ARRAY_INDEX=$AWS_BATCH_JOB_ARRAY_INDEX print-color
    AWS_BATCH_JOB_ARRAY_INDEX=$((AWS_BATCH_JOB_ARRAY_INDEX + 1))
done
```

The following is the output.

```
My favorite color of the rainbow is red.
My favorite color of the rainbow is orange.
My favorite color of the rainbow is yellow.
My favorite color of the rainbow is green.
My favorite color of the rainbow is blue.
My favorite color of the rainbow is indigo.
My favorite color of the rainbow is violet.
```

## **Tutorial: Push your image to Amazon ECR**

Now that you built and tested your Docker container, push it to an image repository. This example uses Amazon ECR, but you can use another registry, such as DockerHub.

Create an Amazon ECR image repository to store your container image. This example only
uses the Amazon CLI, but you can also use the Amazon Web Services Management Console.
For more information, see <u>Creating a Repository</u> in the *Amazon Elastic Container Registry User Guide*.

```
$ aws ecr create-repository --repository-name print-color
```

Tag your print-color image with your Amazon ECR repository URI that was returned from the previous step.

```
$ docker tag print-color aws_account_id.dkr.ecr.region.amazonaws.com/print-color
```

3. Log in to your Amazon ECR registry. For more information, see <u>Registry Authentication</u> in the *Amazon Elastic Container Registry User Guide*.

```
$ aws ecr get-login-password \
    --region region | docker login \
    --username AWS \
    --password-stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

4. Push your image to Amazon ECR.

```
$ docker push aws_account_id.dkr.ecr.region.amazonaws.com/print-color
```

# **Tutorial: Create and register a job definition**

Now that your Docker image is in an image registry, you can specify it in an Amazon Batch job definition. Then, you can use it later to run an array job. This example only uses the Amazon CLI. However, you can also use the Amazon Web Services Management Console. For more information, see <a href="Create a single-node">Create a single-node</a> job definition.

#### To create a job definition

1. Create a file named print-color-job-def.json in your workspace directory and paste the following into it. Replace the image repository URI with your own image's URI.

```
{
  "jobDefinitionName": "print-color",
  "type": "container",
  "containerProperties": {
    "image": "aws_account_id.dkr.ecr.region.amazonaws.com/print-color",
    "resourceRequirements": [
        {
            "type": "MEMORY",
            "value": "250"
        },
        {
            "type": "VCPU",
            "value": "1"
        }
    ]
 }
}
```

2. Register the job definition with Amazon Batch.

```
$ aws batch register-job-definition --cli-input-json file://print-color-job-
def.json
```

# **Tutorial: Submit an Amazon Batch array job**

After you registered your job definition, you can submit an Amazon Batch array job that uses your new container image.

# To submit an Amazon Batch array job

 Create a file named print-color-job.json in your workspace directory and paste the following into it.



### Note

This example uses the job queue mentioned in the the section called "Prerequisites" section.

```
{
  "jobName": "print-color",
  "jobQueue": "existing-job-queue",
  "arrayProperties": {
    "size": 7
  },
  "jobDefinition": "print-color"
}
```

Submit the job to your Amazon Batch job queue. Note the job ID that's returned in the output.

```
$ aws batch submit-job --cli-input-json file://print-color-job.json
```

Describe the job's status and wait for the job to move to SUCCEEDED.

# **Tutorial: View your array job logs**

After your job reaches the SUCCEEDED status, you can view the CloudWatch Logs from the job's container.

#### To view your job's logs in CloudWatch Logs

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. In the left navigation pane, choose **Jobs**.
- 3. For **Job queue**, select a queue.
- 4. In the **Status** section, choose **succeeded**.
- To display all of the child jobs for your array job, select the job ID that was returned in the 5. previous section.
- To see the logs from the job's container, select one of the child jobs and choose View logs.

Filter events					
	Time (UTC +00:00)	Message			
	2018-07-13				
		No older events found at the moment. Retry.			
•	20:16:20	My favorite color of the rainbow is red.			
		No newer events found at the moment. Retry.			

7. View the other child job's logs. Each job returns a different color of the rainbow.

# **Run GPU jobs**

GPU jobs help you to run jobs that use an instance's GPUs.

The following Amazon EC2 GPU-based instance types are supported. For more information, see Amazon EC2 G3 Instances, Amazon EC2 G4 Instances, Amazon EC2 G5 Instances, Amazon EC2 G6 Instances, Amazon EC2 P2 Instances, Amazon EC2 P3 Instances, Amazon EC2 P4d Instances, Amazon EC2 P5 Instances, Amazon EC2 P6 Instances, Amazon EC2 Trn1 Instances, Amazon EC2 Trn2 Instances, Amazon EC2 Inf1 Instances, Amazon EC2 Inf2 Instances, Amazon EC2 Dl1 Instances, and Amazon EC2 Dl2 Instances.

Instance type	GPUs	GPU memory	vCPUs	Memory	Network bandwidth
g3s.xlarge	1	8 GiB	4	30.5 GiB	10 Gbps
g3.4xlarge	1	8 GiB	16	122 GiB	Up to 10 Gbps
g3.8xlarge	2	16 GiB	32	244 GiB	10 Gbps
g3.16xlarge	4	32 GiB	64	488 GiB	25 Gbps
g4dn.xlarge	1	16 GiB	4	16 GiB	Up to 25 Gbps
g4dn.2xlarge	1	16 GiB	8	32 GiB	Up to 25 Gbps
g4dn.4xlarge	1	16 GiB	16	64 GiB	Up to 25 Gbps

Instance type	GPUs	GPU memory	vCPUs	Memory	Network bandwidth
g4dn.8xlarge	1	16 GiB	32	128 GiB	50 Gbps
g4dn.12xlarge	4	64 GiB	48	192 GiB	50 Gbps
g4dn.16xlarge	1	16 GiB	64	256 GiB	50 Gbps
g5.xlarge	1	24 GiB	4	16 GiB	Up to 10 Gbps
g5.2xlarge	1	24 GiB	8	32 GiB	Up to 10 Gbps
g5.4xlarge	1	24 GiB	16	64 GiB	Up to 25 Gbps
g5.8xlarge	1	24 GiB	32	128 GiB	25 Gbps
g5.16xlarge	1	24 GiB	64	256 GiB	25 Gbps
g5.12xlarge	4	96 GiB	48	192 GiB	40 Gbps
g5.24xlarge	4	96 GiB	96	384 GiB	50 Gbps
g5.48xlarge	8	192 GiB	192	768 GiB	100 Gbps
g5g.xlarge	1	16 GB	4	8 GB	Up to 10 Gbps
g5g.2xlarge	1	16 GB	8	16 GB	Up to 10 Gbps
g5g.4xlarge	1	16 GB	16	32 GB	Up to 10 Gbps
g5g.8xlarge	1	16 GB	32	64 GB	12 Gbps
g5g.16xlarge	2	32 GB	64	128 GB	25 Gbps
g5g.metal	2	32 GB	64	128 GB	25 Gbps
g6.xlarge	1	24 GiB	4	16 GiB	Up to 10 Gbps
g6.2xlarge	1	24 GiB	8	32 GiB	Up to 10 Gbps
g6.4xlarge	1	24 GiB	16	64 GiB	Up to 25 Gbps

Instance type	GPUs	GPU memory	vCPUs	Memory	Network bandwidth
g6.8xlarge	1	24 GiB	32	128 GiB	25 Gbps
g6.16xlarge	1	24 GiB	64	256 GiB	25 Gbps
g6.12xlarge	4	96 GiB	48	192 GiB	40 Gbps
g6.24xlarge	4	96 GiB	96	384 GiB	50 Gbps
g6.48xlarge	8	192 GiB	192	768 GiB	100 Gbps
g6e.xlarge	1	48 GB	4	32 GB	Up to 20 Gbps
g6e.2xlarge	1	48 GB	8	64 GB	Up to 20 Gbps
g6e.4xlarge	1	48 GB	16	128 GB	20 Gbps
g6e.8xlarge	1	48 GB	32	256 GB	25 Gbps
g6e.16xlarge	1	48 GB	64	512 GB	35 Gbps
g6e.12xlarge	4	192 GB	48	384 GB	100 Gbps
g6e.24xlarge	4	192 GB	96	768 GB	200 Gbps
g6e.48xlarge	8	384 GB	192	1536 GB	400 Gbps
gr6.4xlarge	1	24 GB	16	128 GB	Up to 25 Gbps
gr6.8xlarge	1	24 GB	32	256 GB	25 Gbps
p2.xlarge	1	12 GiB	4	61 GiB	High
p2.8xlarge	8	96 GiB	32	488 GiB	10 Gbps
p2.16xlarge	16	192 GiB	64	732 GiB	20 Gbps
p3.2xlarge	1	16 GiB	8	61 GiB	Up to 10 Gbps
p3.8xlarge	4	64 GiB	32	244 GiB	10 Gbps

Instance type	GPUs	GPU memory	vCPUs	Memory	Network bandwidth
p3.16xlarge	8	128 GiB	64	488 GiB	25 Gbps
p3dn.24xlarge	8	256 GiB	96	768 GiB	100 Gbps
p4d.24xlarge	8	320 GB	96	1152 GB	400 Gbps
p4de.24xlarge	8	640 GB	96	1152 GB	400 Gbps
p5.48xlarge	8	640 GiB	192	2 TiB	3200 Gbps
p5e.48xlarge	8	1128 GB	192	2 TB	3200 Gbps
p5en.48xlarge	8	1128 GB	192	2 TB	3200 Gbps
p6-b200.4 8xlarge	8	1440 GB	192	2 TB	100 Gbps
trn1.2xlarge	1	32 GB	8	32 GB	Up to 12.5 Gbps
trn1.32xlarge	16	512 GB	128	512 GB	800 Gbps
trn1n.32xlarge	16	512 GB	128	512 GB	1600 Gbps
trn2.48xlarge	16	1.5 TB	192	2 TB	3.2 Tbps
inf1.xlarge	1	8 GB	4	8 GB	Up to 25 Gbps
inf1.2xlarge	1	8 GB	8	16 GB	Up to 25 Gbps
inf1.6xlarge	4	32 GB	24	48 GB	25 Gbps
inf1.24xlarge	16	128 GB	96	192 GB	100 Gbps
inf2.xlarge	1	32 GB	4	16 GB	Up to 15 Gbps
inf2.8xlarge	1	32 GB	32	128 GB	Up to 25 Gbps
inf2.24xlarge	6	192 GB	96	384 GB	50 Gbps

Instance type	GPUs	GPU memory	vCPUs	Memory	Network bandwidth
inf2.48xlarge	12	384 GB	192	768 GB	100 Gbps
dl1.24xlarge	8	256 GB	96	768 GB	400 Gbps
dl2q.24xlarge	8	128 GB	96	768 GB	100 Gbps

### Note

For GPU jobs Amazon Batch only supports instance types that have NVIDIA GPUs. For example, the <u>G4ad</u> family is not supported for GPU scheduling. You can still use <u>G4ad</u> on Amazon Batch by defining only the vcpu and memory requirements in the job definition, then accessing the host GPUs directly through customization in an Amazon EC2 <u>launch</u> <u>template user data</u> with an Amazon ECS or Amazon EKS compute optimized AMI, or a customized AMI for using AMD GPUs.

Instance types that use an ARM64 architecture are supported for GPU jobs on custom AMIs provided to Amazon Batch or Amazon EC2 user data to access the GPUs by customized code and configurations. For example, the G5g instance family.

The <u>resourceRequirements</u> parameter for the job definition specifies the number of GPUs to be pinned to the container. This number of GPUs isn't available to any other job that runs on that instance for the duration of that job. All instance types in a compute environment that run GPU jobs must be from the p3, p4, p5, p6, g3, g3s, g4, g5, or g6 instance families. If this isn't done a GPU job might get stuck in the RUNNABLE status.

Jobs that don't use the GPUs can be run on GPU instances. However, they might cost more to run on the GPU instances than on similar non-GPU instances. Depending on the specific vCPU, memory, and time needed, these non-GPU jobs might block GPU jobs from running.

#### **Topics**

- Tutorial: Create a GPU-based Kubernetes cluster on Amazon EKS
- Tutorial: Create an Amazon EKS GPU job definition
- Tutorial: Run a GPU job in your Amazon EKS cluster

### Tutorial: Create a GPU-based Kubernetes cluster on Amazon EKS

Before you create a GPU-based Kubernetes cluster on Amazon EKS, you must have completed the steps in <u>Getting started with Amazon Batch on Amazon EKS</u>. In addition, also consider the following:

- Amazon Batch supports instance types with NVIDIA GPUs.
- By default, Amazon Batch selects the Amazon EKS accelerated AMI with the Kubernetes version that matches your Amazon EKS cluster control plane version.

```
$ cat <<EOF > ./batch-eks-gpu-ce.json
{
  "computeEnvironmentName": "My-Eks-GPU-CE1",
  "type": "MANAGED",
  "state": "ENABLED",
  "eksConfiguration": {
    "eksClusterArn": "arn:aws-cn:eks:<region>:<account>:cluster/<cluster-name>",
    "kubernetesNamespace": "my-aws-batch-namespace"
  },
  "computeResources": {
    "type": "EC2",
    "allocationStrategy": "BEST_FIT_PROGRESSIVE",
    "minvCpus": 0,
    "maxvCpus": 1024,
    "instanceTypes": [
      "p3dn.24xlarge",
      "p4d.24xlarge"
    ],
    "subnets": [
        "<eks-cluster-subnets-with-access-to-internet-for-image-pull>"
    ],
    "securityGroupIds": [
        "<eks-cluster-sg>"
    ],
    "instanceRole": "<eks-instance-profile>"
  }
}
EOF
$ aws batch create-compute-environment --cli-input-json file://./batch-eks-gpu-ce.json
```

Amazon Batch doesn't manage the NVIDIA GPU device plugin on your behalf. You must install this plugin into your Amazon EKS cluster and allow it to target the Amazon Batch nodes. For more information, see Enabling GPU Support in Kubernetes on GitHub.

To configure the NVIDIA device plugin (DaemonSet) to target the Amazon Batch nodes, run the following commands.

```
# pull nvidia daemonset spec
$ curl -0 https://raw.githubusercontent.com/NVIDIA/k8s-device-plugin/v0.12.2/nvidia-
device-plugin.yml
# using your favorite editor, add Batch node toleration
# this will allow the DaemonSet to run on Batch nodes
- key: "batch.amazonaws.com/batch-node"
    operator: "Exists"

$ kubectl apply -f nvidia-device-plugin.yml
```

We do not recommend that you mix compute-based (CPU and memory) workloads with GPU-based workloads in the same pairings of compute environment and job queue. This is because compute jobs can use up GPU capacity.

To attach job queues, run the following commands.

```
$ cat <<EOF > ./batch-eks-gpu-jq.json
{
    "jobQueueName": "My-Eks-GPU-JQ1",
    "priority": 10,
    "computeEnvironmentOrder": [
        {
            "order": 1,
            "computeEnvironment": "My-Eks-GPU-CE1"
        }
    ]
}
EOF

$ aws batch create-job-queue --cli-input-json file://./batch-eks-gpu-jq.json
```

# Tutorial: Create an Amazon EKS GPU job definition

Only nvidia.com/gpu is supported at this time and resource value that you set must be a whole number. You can't use fractions of GPU. For more information, see <a href="Schedule GPUs">Schedule GPUs</a> in the Kubernetes documentation.

To register a GPU job definition for Amazon EKS, run the following commands.

```
$ cat <<EOF > ./batch-eks-gpu-jd.json
}
    "jobDefinitionName": "MyGPUJobOnEks_Smi",
    "type": "container",
    "eksProperties": {
        "podProperties": {
            "hostNetwork": true,
            "containers": [
                    "image": "nvcr.io/nvidia/cuda:10.2-runtime-centos7",
                    "command": ["nvidia-smi"],
                    "resources": {
                         "limits": {
                             "cpu": "1",
                             "memory": "1024Mi",
                             "nvidia.com/gpu": "1"
                         }
                    }
                }
            ]
        }
    }
}
EOF
$ aws batch register-job-definition --cli-input-json file://./batch-eks-gpu-jd.json
```

# Tutorial: Run a GPU job in your Amazon EKS cluster

The GPU resource is non-compressible. Amazon Batch creates a pod spec for GPU jobs where the value of **request** equals the value of **limits**. This is a Kubernetes requirement.

To submit a GPU job, run the following commands.

# Search Amazon Batch jobs in a job queue

You can search and filter your jobs in Amazon Batch using Job search. This feature provides an option to search within a pre-existing job queue and filter its jobs by one of five options (Created after, Created before, Job name, Job definition, and Status). For the process to use this feature, complete the Tutorial: Search and filter Amazon Batch jobs procedure.

To search using multiple criteria simultaneously, use the Beta Job search feature. For example, you can include any or all of the following filters: Job status, Date range, and Additional criteria (such as, a job name, job definition, or job ID). For the process to use this feature, see the <u>Tutorial</u>: (Beta) Search and filter Amazon Batch jobs procedure.

#### **Topics**

- Tutorial: Search and filter Amazon Batch jobs
- Tutorial: (Beta) Search and filter Amazon Batch jobs

# **Tutorial: Search and filter Amazon Batch jobs**

Use this procedure to list all the jobs in a job queue in the Amazon Batch console. Optionally, use the Filter type dropdown list to narrow the results based on the criteria you specify.

- 1. Navigate to the Amazon Batch console.
- 2. In the Navigation pane, choose **Jobs**.
- 3. Expand the **Job queue** dropdown list and choose the job queue that you want to search within.



#### Note

You can search for jobs within only one job queue at a time.

- 4. (Optional) Choose the filter type you want to apply.
  - Created after: Returns jobs created after a specific date and time (in UTC). If you select this option, two filter value fields appear where you can specify the earliest date and time that the search will return.
  - Created before: Returns jobs created before a specific date and time (in UTC). If you select this option, two filter value fields appear where you can specify the latest date and time that the search will return.
  - Job definition: Returns jobs created with a specific job definition name. If you select this option, a field appears where you can specify a job's complete or partial definition (if you add an asterisk (\*) at the end, the filter matches any job definition that begins with the string before the asterisk).
  - Job name: Returns jobs that have a specific job name. If you select this option, a field appears where you can specify a job's complete or partial name (if you add an asterisk (\*) at the end, the filter matches any job name that begins with the string before the asterisk).
  - **Status**: Returns jobs that have the selected job status. If you select this option, a dropdown list appears for you to expand and select a specific status. For more information, see Job states.
- (Optional) Choose Load all jobs to return all jobs within a job queue in the last seven days and without the previous filters. If you choose this option, the following message appears:



#### Note

Amazon Batch now stores seven (7) days of job history. Choosing to load all jobs will increase the time it takes to access your jobs and start searching and filtering. Using a filter type instead will provide more immediate results. Are you sure you want to load all jobs.

A Choose **Confirm** to close the message.

• If you choose this option, a dropdown list appears where you can choose similar filter options: **Status**, **Job name**, and **Job ID**. If you choose one of these options, you can further filter your search criteria where the specified filter is equal to, is not equal to, contains, or does not contain the criteria you enter.

When the options are complete, choose **Search**.

# **Tutorial: (Beta) Search and filter Amazon Batch jobs**

Use this procedure to list all the jobs in a job queue in the Amazon Batch console and using the beta search features. Optionally, use the Search options section to narrow the results based on the criteria you specify. This beta search feature allows you to combine search criteria to narrow the results. However, you can define only one criteria per filter.

- 1. Navigate to the Amazon Batch console.
- 2. In the Navigation pane, choose **Jobs**.
- 3. Ensure that the **Search jobs - beta** option is active.
- Expand the Job queue dropdown list and choose the job queue that you want to search within. 4.



#### Note

You can search for jobs within only one job queue at a time.

- (Optional) Expand the **Job status** dropdown list and select one or more job states. For more information, see Job states.
- (Optional) Choose Date range to filter the results based on a date and time range.
  - Choose Relative mode to search for jobs that ran within a time range counting backwards from the current date and time.
  - Choose Absolute mode to search for jobs that ran within a date and time range that you specify.
- (Optional) In the Additional criteria field, enter keywords to include in the search results. For example, you can use this field to search by job name, job definition, or job ID. Depending on the property, there may be additional operators, such as equals (=) or starts with (^) that you must define.
- When the options are complete, choose **Search**. 8.



You can choose the **settings** icon to customize column that appear and their order in the Search results table.

# **Networking modes for Amazon Batch jobs**

The following table describes the networking modes and typical usage for Amazon Batch job types. See the links in the "Job type" column for more details regarding considerations and behaviors.

Job Type	Supported Network Mode(s)	Typical Usage
ECS-EC2 simple job	host	Used for the highest scalable embarrassingly parallel batch workloads that only require egress to vpc defined in Compute Environment.
ECS-EC2 multi-node parallel job	awsvpc	Used for tightly-coupled, multi-host (node) distribut ed workloads modeled as a single job with coordinated communications between task nodes.
ECS-Fargate simple job	awsvpc	True serverless for embarrass ingly parallel batch workloads . Typically lowest TCO and highest container isolation job model.
EKS-EC2 simple job	host and pod	Used for highly scalable embarrassingly parallel batch workloads that only require egress to vpc defined

Job Type	Supported Network Mode(s)	Typical Usage
		in Compute Environment.  Default is host networking.
EKS-EC2 multi-node parallel job	host and pod	Used for tightly-coupled, multi-host (node) distribut ed workloads modeled as a single job with coordinated communications between pod nodes. Default is host networking.

# Tutorial: Send Amazon Batch job logs to CloudWatch Logs

You can configure your Amazon Batch jobs to send log information to Amazon CloudWatch Logs. This way, you can view different logs from your jobs in one convenient location. For more information, see Using CloudWatch Logs with Amazon Batch.

You can also use **Job logs** in the Amazon Batch console to monitor or troubleshoot an Amazon Batch job.

- Open the Amazon Batch console. 1.
- Choose Jobs. 2.
- For **Job queue**, choose the job queue that you want.



If there are several jobs in the job queue, you can turn on Searching and filtering to find a job faster. For more information, see Search Amazon Batch jobs in a job queue.

- 4. For **Status**, choose the job status that you want.
- 5. Choose the job that you want.
- 6. On the **Details** page, scroll down to **Job Logs**.
- 7. Choose **Retrieve logs**.

For **Authorization required**, enter **OK**, and then choose **Authorize** to accept Amazon CloudWatch charges.



#### Note

To revoke your authorization for CloudWatch charges:

- 1. In the left navigation pane, choose **Permissions**.
- 2. For **Job logs**, choose **Edit**.
- 3. Clear the Authorize Batch to use CloudWatch check box.
- 4. Choose **Save changes**.
- Review the log data for the Amazon Batch job.



#### (i) Tip

You can filter the log based on **Keywords**, **Max results**, and **Sorting**. You can also choose one of the default time intervals or create a custom interval to customize the results.

# **Tutorial: Review Amazon Batch job information**

You can review Amazon Batch job information such as status, job definition and container information.

- 1. Open the Amazon Batch console.
- 2. Choose Jobs.
- For **Job queue**, choose the job queue that you want.



If there are several jobs in the job queue, you can turn on Search and filter to find a job faster. For more information, see Search Amazon Batch jobs in a job queue.

Choose the job that you want. 4.



## Note

You can also use the Amazon Command Line Interface (Amazon CLI) to view details about an Amazon Batch job. For more information, see describe-jobs in the Amazon CLI Command Reference.

# **Security in Amazon Batch**

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security sensitive organizations.

Security is a shared responsibility between Amazon and you. The <u>shared responsibility model</u> describes this as security *of* the cloud and security *in* the cloud:

- Security of the cloud Amazon is responsible for protecting the infrastructure that runs
   Amazon services in the Amazon Web Services Cloud. Amazon also provides you with services
   that you can use securely. Third-party auditors regularly test and verify the effectiveness of our
   security as part of the <u>Amazon Compliance Programs</u>. To learn about the compliance programs
   that apply to Amazon Batch, see <u>Amazon Services in Scope by Compliance Program</u>.
- Security in the cloud Your responsibility is determined by the Amazon service that you use.
   You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Batch. The following topics show you how to configure Amazon Batch to meet your security and compliance objectives. You also learn how to use other Amazon services that help you to monitor and secure your Amazon Batch resources.

#### **Topics**

- Identity and Access Management for Amazon Batch
- Amazon Batch IAM policies, roles, and permissions
- Amazon Batch IAM execution role
- Create a virtual private cloud
- Use an interface endpoint to Access Amazon Batch
- Compliance validation for Amazon Batch
- Infrastructure security in Amazon Batch
- Cross-service confused deputy prevention
- Logging Amazon Batch API calls with Amazon CloudTrail
- Troubleshoot Amazon Batch IAM

# **Identity and Access Management for Amazon Batch**

Amazon Identity and Access Management (IAM) is an Amazon Web Services service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Batch resources. IAM is an Amazon Web Services service that you can use with no additional charge.

#### **Topics**

- Audience
- · Authenticating with identities
- Managing access using policies
- How Amazon Batch works with IAM
- Identity-based policy examples for Amazon Batch
- Amazon managed policies for Amazon Batch

#### **Audience**

How you use Amazon Identity and Access Management (IAM) differs, depending on the work that you do in Amazon Batch.

**Service user** – If you use the Amazon Batch service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon Batch features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon Batch, see Troubleshoot Amazon Batch IAM.

**Service administrator** – If you're in charge of Amazon Batch resources at your company, you probably have full access to Amazon Batch. It's your job to determine which Amazon Batch features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Batch, see <a href="How Amazon Batch works with IAM">How Amazon Batch works with IAM</a>.

**IAM administrator** – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon Batch. To view example Amazon Batch identity-based policies that you can use in IAM, see Identity-based policy examples for Amazon Batch.

# **Authenticating with identities**

Authentication is how you sign in to Amazon using your identity credentials. You must be *authenticated* (signed in to Amazon) as the Amazon Web Services account root user, as an IAM user, or by assuming an IAM role.

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see <a href="Amazon Signature Version 4">Amazon April requests</a> in the IAM User Guide.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <a href="Amazon Multi-factor authentication"><u>Amazon Multi-factor authentication in IAM in the IAM User Guide.</u></a>

#### Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services services and resources in the account. This identity is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see <u>Tasks that require root user credentials</u> in the *IAM User Guide*.

# **Federated identity**

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access Amazon Web Services services by using temporary credentials.

A federated identity is a user from your enterprise user directory, a web identity provider, the Amazon Directory Service, or any user that accesses Amazon Web Services services by using credentials provided through an identity source. When federated identities access Amazon Web Services accounts, they assume roles, and the roles provide temporary credentials.

Authenticating with identities

## IAM users and groups

An <u>IAM user</u> is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for use cases that require long-term credentials</u> in the <u>IAM User Guide</u>.

An <u>IAM group</u> is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see <u>Use cases for IAM users</u> in the *IAM User Guide*.

#### IAM roles

An <u>IAM role</u> is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the Amazon Web Services Management Console, you can <u>switch from a user to an IAM role (console)</u>. You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see <u>Methods to assume a role in the IAM User Guide</u>.

IAM roles with temporary credentials are useful in the following situations:

- Federated user access To assign permissions to a federated identity, you create a role
  and define permissions for the role. When a federated identity authenticates, the identity
  is associated with the role and is granted the permissions that are defined by the role. For
  information about roles for federation, see <a href="Create a role for a third-party identity provider">Create a role for a third-party identity provider</a>
  (federation) in the IAM User Guide.
- **Temporary IAM user permissions** An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- Cross-account access You can use an IAM role to allow someone (a trusted principal) in a
  different account to access resources in your account. Roles are the primary way to grant cross-

Authenticating with identities 337

account access. However, with some Amazon Web Services services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see <a href="Cross account resource access in IAM">Cross account resource access in IAM</a> in the IAM User Guide.

- Cross-service access Some Amazon Web Services services use features in other Amazon Web Services services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
  - Forward access sessions (FAS) When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see <a href="Forward access sessions">Forward access sessions</a>.
  - Service role A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM.
     For more information, see <u>Create a role to delegate permissions to an Amazon Web Services</u> service in the *IAM User Guide*.
  - Service-linked role A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- Applications running on Amazon EC2 You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see <a href="Use an IAM role to grant permissions to applications running on Amazon EC2 instances">Use an IAM role to grant permissions to applications running on Amazon EC2 instances</a> in the IAM User Guide.

# Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the structure and contents of JSON policy documents, see <a href="Overview of JSON policies">Overview of JSON policies</a> in the IAM User Guide.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam: GetRole action. A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

# **Identity-based policies**

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <u>Define custom IAM permissions with customer managed policies</u> in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see <a href="Choose between managed policies and inline policies">Choose between managed policies and inline policies</a> in the *IAM User Guide*.

# **Resource-based policies**

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that

support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must specify a principal in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

### Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, Amazon WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

## Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- Permissions boundaries A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see Permissions boundaries for IAM entities in the IAM User Guide.
- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see <a href="Service control policies">Service control policies</a> in the Amazon Organizations User Guide.

• Resource control policies (RCPs) – RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the Amazon Web Services account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of Amazon Web Services services that support RCPs, see Resource control policies (RCPs) in the Amazon Organizations User Guide.

• Session policies – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see Session policies in the IAM User Guide.

# Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

### How Amazon Batch works with IAM

Before you use IAM to manage access to Amazon Batch, learn what IAM features are available to use with Amazon Batch.

#### IAM features you can use with Amazon Batch

IAM feature	Amazon Batch support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	Yes
ACLs	No

IAM feature	Amazon Batch support
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Principal permissions	Yes
Service roles	Yes
Service-linked roles	Yes

To get a high-level view of how Amazon Batch and other Amazon services work with most IAM features, see Amazon services that work with IAM in the IAM User Guide.

# **Identity-based policies for Amazon Batch**

### Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <u>Define custom IAM permissions with customer managed policies</u> in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see <a href="IAM JSON policy elements reference">IAM JSON policy elements reference</a> in the IAM User Guide.

### **Identity-based policy examples for Amazon Batch**

To view examples of Amazon Batch identity-based policies, see <u>Identity-based policy examples for</u> Amazon Batch.

# **Policy actions for Amazon Batch**

# Supports policy actions: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated Amazon API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon Batch actions, see <u>Actions Defined by Amazon Batch</u> in the *Service Authorization Reference*.

Policy actions in Amazon Batch use the following prefix before the action:

```
batch
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
    "batch:action1",
    "batch:action2"
]
```

You can specify multiple actions using wildcards (\*). For example, to specify all actions that begin with the word Describe, include the following action:

```
"Action": "batch:Describe*"
```

To view examples of Amazon Batch identity-based policies, see <u>Identity-based policy examples for</u> Amazon Batch.

# **Policy resources for Amazon Batch**

# Supports policy resources: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its <a href="Management-Amazon Resource Name">Amazon Resource Name</a> (ARN). You can do this for actions that support a specific resource type, known as resource-level permissions.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (\*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To see a list of Amazon Batch resource types and their ARNs, see <u>Resources Defined by Amazon</u>

<u>Batch</u> in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see <u>Actions Defined by Amazon Batch</u>.

### **Policy condition keys for Amazon Batch**

#### Supports service-specific policy condition keys: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use <u>condition operators</u>, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, Amazon evaluates them using a logical AND operation. If you specify multiple values for a single condition key, Amazon evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the IAM User Guide.

Amazon supports global condition keys and service-specific condition keys. To see all Amazon global condition keys, see Amazon global condition context keys in the *IAM User Guide*.

To see a list of Amazon Batch condition keys, see <u>Condition Keys for Amazon Batch</u> in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see Actions Defined by Amazon Batch.

# Attribute-based access control (ABAC) with Amazon Batch

Supports ABAC (tags in policies): Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In Amazon, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many Amazon resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the <u>condition element</u> of a policy using the aws:ResourceTag/key-name, aws:RequestTag/key-name, or aws:TagKeys condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see <u>Define permissions with ABAC authorization</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>Use attribute-based access control</u> (ABAC) in the *IAM User Guide*.

### Use temporary credentials with Amazon Batch

### Supports temporary credentials: Yes

Some Amazon Web Services services don't work when you sign in using temporary credentials. For additional information, including which Amazon Web Services services work with temporary credentials, see Amazon Web Services services that work with IAM in the IAM User Guide.

You are using temporary credentials if you sign in to the Amazon Web Services Management Console using any method except a user name and password. For example, when you access Amazon using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see <a href="Switch from a user to an IAM role (console)">Switch from a user to an IAM role (console)</a> in the IAM User Guide.

You can manually create temporary credentials using the Amazon CLI or Amazon API. You can then use those temporary credentials to access Amazon. Amazon recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see Temporary security credentials in IAM.

### **Cross-service principal permissions for Amazon Batch**

#### **Supports forward access sessions (FAS):** Yes

When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

#### Service roles for Amazon Batch

#### Supports service roles: Yes

A service role is an IAM role that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see Create a role to delegate permissions to an Amazon Web Services service in the IAM User Guide.

#### Marning

Changing the permissions for a service role might break Amazon Batch functionality. Edit service roles only when Amazon Batch provides guidance to do so.

#### Service-linked roles for Amazon Batch

### Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see Amazon services that work with IAM. Find a service in the table that includes a Yes in the Service-linked role column. Choose the **Yes** link to view the service-linked role documentation for that service.

# **Identity-based policy examples for Amazon Batch**

By default, users and roles don't have permission to create or modify Amazon Batch resources. They also can't perform tasks by using the Amazon Web Services Management Console, Amazon Command Line Interface (Amazon CLI), or Amazon API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see Create IAM policies (console) in the IAM User Guide.

For details about actions and resource types defined by Amazon Batch, including the format of the ARNs for each of the resource types, see <u>Actions, Resources, and Condition Keys for Amazon Batch</u> in the *Service Authorization Reference*.

#### **Topics**

- Policy best practices
- Using the Amazon Batch console
- Allow users to view their own permissions

# **Policy best practices**

Identity-based policies determine whether someone can create, access, or delete Amazon Batch resources in your account. These actions can incur costs for your Amazon Web Services account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- Get started with Amazon managed policies and move toward least-privilege permissions
  - To get started granting permissions to your users and workloads, use the *Amazon managed policies* that grant permissions for many common use cases. They are available in your Amazon Web Services account. We recommend that you reduce permissions further by defining Amazon customer managed policies that are specific to your use cases. For more information, see <u>Amazon managed policies</u> or <u>Amazon managed policies</u> for job functions in the *IAM User Guide*.
- Apply least-privilege permissions When you set permissions with IAM policies, grant only the
  permissions required to perform a task. You do this by defining the actions that can be taken on
  specific resources under specific conditions, also known as least-privilege permissions. For more
  information about using IAM to apply permissions, see <a href="Policies and permissions in IAM">Policies and permissions in IAM</a> in the
  IAM User Guide.

Use conditions in IAM policies to further restrict access – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific Amazon Web Services service, such as Amazon CloudFormation. For more information, see <a href="IAM JSON policy elements: Condition">IAM User Guide</a>.

- Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see <u>Validate policies with IAM Access Analyzer</u> in the *IAM User Guide*.
- Require multi-factor authentication (MFA) If you have a scenario that requires IAM users or a
  root user in your Amazon Web Services account, turn on MFA for additional security. To require
  MFA when API operations are called, add MFA conditions to your policies. For more information,
  see Secure API access with MFA in the IAM User Guide.

For more information about best practices in IAM, see <u>Security best practices in IAM</u> in the *IAM User Guide*.

# Using the Amazon Batch console

To access the Amazon Batch console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon Batch resources in your Amazon Web Services account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the Amazon CLI or the Amazon API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Amazon Batch console, also attach the Amazon Batch ConsoleAccess or ReadOnly Amazon managed policy to the entities. For more information, see Adding permissions to a user in the *IAM User Guide*.

# Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the Amazon CLI or Amazon API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws-cn:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
        }
    ]
}
```

# **Amazon managed policies for Amazon Batch**

You can use Amazon managed policies for simpler identity access management for your team and provisioned Amazon resources. Amazon managed policies cover a variety of common use cases, are available by default in your Amazon account, and are maintained and updated on your behalf. You can't change the permissions in Amazon managed policies. If you require greater flexibility, you can alternatively choose to create IAM customer managed policies. This way, you can provide your team provisioned resources with only the exact permissions they need.

For more information about Amazon managed policies, see <u>Amazon managed policies</u> in the *IAM User Guide*.

Amazon services maintain and update Amazon managed policies on your behalf. Periodically, Amazon services add additional permissions to an Amazon managed policy. Amazon managed policies are most likely updated when a new feature launch or operation becomes available. These updates automatically affect all identities (users, groups, and roles) where the policy is attached. However, they don't remove permissions or break your existing permissions.

Additionally, Amazon supports managed policies for job functions that span multiple services. For example, the ReadOnlyAccess Amazon managed policy provides read-only access to all Amazon services and resources. When a service launches a new feature, Amazon adds read-only permissions for new operations and resources. For a list and descriptions of job function policies, see <a href="Amazon managed policies">Amazon managed policies for job functions in the IAM User Guide</a>.

# Amazon managed policy: BatchServiceRolePolicy

The **BatchServiceRolePolicy** managed IAM policy is used by the <u>AWSServiceRoleForBatch</u> service-linked role. This allows Amazon Batch to perform actions on your behalf. You can't attach this policy to your IAM entities. For more information, see <u>Use service-linked roles for Amazon</u> Batch.

This policy allows Amazon Batch to complete the following actions on specific resources:

- autoscaling Allows Amazon Batch to create and manage Amazon EC2 Auto Scaling resources. Amazon Batch creates and manages Amazon EC2 Auto Scaling groups for most compute environments.
- ec2 Allows Amazon Batch to control the lifecycle of Amazon EC2 instances as well as create and manage launch templates and tags. Amazon Batch creates and manages EC2 Spot Fleet requests for some EC2 Spot compute environments.

• ecs - Allows Amazon Batch to create and managed Amazon ECS clusters, task definitions and tasks for job execution.

- eks Allows Amazon Batch to describe the Amazon EKS cluster resource for validations.
- iam Allows Amazon Batch to validate and pass roles provided by owner to Amazon EC2,
   Amazon EC2 Auto Scaling and Amazon ECS.
- logs Allows Amazon Batch to create and manage log groups and log streams for Amazon Batch jobs.

**JSON** 

```
}
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AWSBatchPolicyStatement1",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeAccountAttributes",
                "ec2:DescribeInstances",
                "ec2:DescribeInstanceStatus",
                "ec2:DescribeInstanceAttribute",
                "ec2:DescribeSubnets",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeKeyPairs",
                "ec2:DescribeImages",
                "ec2:DescribeImageAttribute",
                "ec2:DescribeSpotInstanceRequests",
                "ec2:DescribeSpotFleetInstances",
                "ec2:DescribeSpotFleetRequests",
                "ec2:DescribeSpotPriceHistory",
                "ec2:DescribeSpotFleetRequestHistory",
                "ec2:DescribeVpcClassicLink",
                "ec2:DescribeLaunchTemplateVersions",
                "ec2:RequestSpotFleet",
                "autoscaling:DescribeAccountLimits",
                "autoscaling:DescribeAutoScalingGroups",
                "autoscaling:DescribeLaunchConfigurations",
                "autoscaling:DescribeAutoScalingInstances",
                "autoscaling:DescribeScalingActivities",
                "eks:DescribeCluster",
```

```
"ecs:DescribeClusters",
                "ecs:DescribeContainerInstances",
                "ecs:DescribeTaskDefinition",
                "ecs:DescribeTasks",
                "ecs:ListClusters",
                "ecs:ListContainerInstances",
                "ecs:ListTaskDefinitionFamilies",
                "ecs:ListTaskDefinitions",
                "ecs:ListTasks",
                "ecs:DeregisterTaskDefinition",
                "ecs:TagResource",
                "ecs:ListAccountSettings",
                "logs:DescribeLogGroups",
                "iam:GetInstanceProfile",
                "iam:GetRole"
            ],
            "Resource": "*"
        },
        {
            "Sid": "AWSBatchPolicyStatement2",
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogGroup",
                "logs:CreateLogStream"
            ],
            "Resource": "arn:aws-cn:logs:*:*:log-group:/aws/batch/job*"
        },
            "Sid": "AWSBatchPolicyStatement3",
            "Effect": "Allow",
            "Action": [
                "logs:PutLogEvents"
            ],
            "Resource": "arn:aws-cn:logs:*:*:log-group:/aws/batch/job*:log-
stream:*"
        },
        }
            "Sid": "AWSBatchPolicyStatement4",
            "Effect": "Allow",
            "Action": [
                "autoscaling:CreateOrUpdateTags"
            "Resource": "*",
            "Condition": {
```

```
"Null": {
            "aws:RequestTag/AWSBatchServiceTag": "false"
        }
    }
},
{
    "Sid": "AWSBatchPolicyStatement5",
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": [
        11 * 11
    ],
    "Condition": {
        "StringEquals": {
            "iam:PassedToService": [
                "ec2.amazonaws.com",
                "ec2.amazonaws.com.cn",
                "ecs-tasks.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AWSBatchPolicyStatement6",
    "Effect": "Allow",
    "Action": "iam:CreateServiceLinkedRole",
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": [
                "spot.amazonaws.com",
                "spotfleet.amazonaws.com",
                "autoscaling.amazonaws.com",
                "ecs.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AWSBatchPolicyStatement7",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateLaunchTemplate"
    ],
```

```
"Resource": "*",
            "Condition": {
                "Null": {
                    "aws:RequestTag/AWSBatchServiceTag": "false"
                }
            }
        },
        {
            "Sid": "AWSBatchPolicyStatement8",
            "Effect": "Allow",
            "Action": [
                "ec2:TerminateInstances",
                "ec2:CancelSpotFleetRequests",
                "ec2:ModifySpotFleetRequest",
                "ec2:DeleteLaunchTemplate"
            ],
            "Resource": "*",
            "Condition": {
                "Null": {
                    "aws:ResourceTag/AWSBatchServiceTag": "false"
                }
            }
        },
        {
            "Sid": "AWSBatchPolicyStatement9",
            "Effect": "Allow",
            "Action": [
                "autoscaling:CreateLaunchConfiguration",
                "autoscaling:DeleteLaunchConfiguration"
            ],
            "Resource": "arn:aws-
cn:autoscaling:*:*:launchConfiguration:*:launchConfigurationName/AWSBatch*"
        },
        }
            "Sid": "AWSBatchPolicyStatement10",
            "Effect": "Allow",
            "Action": [
                "autoscaling:CreateAutoScalingGroup",
                "autoscaling:UpdateAutoScalingGroup",
                "autoscaling:SetDesiredCapacity",
                "autoscaling:DeleteAutoScalingGroup",
                "autoscaling:SuspendProcesses",
                "autoscaling:PutNotificationConfiguration",
                "autoscaling:TerminateInstanceInAutoScalingGroup"
```

```
],
            "Resource": "arn:aws-
cn:autoscaling:*:*:autoScalingGroup:*:autoScalingGroupName/AWSBatch*"
        {
            "Sid": "AWSBatchPolicyStatement11",
            "Effect": "Allow",
            "Action": [
                "ecs:DeleteCluster",
                "ecs:DeregisterContainerInstance",
                "ecs:RunTask",
                "ecs:StartTask",
                "ecs:StopTask"
            ],
            "Resource": "arn:aws-cn:ecs:*:*:cluster/AWSBatch*"
        },
        {
            "Sid": "AWSBatchPolicyStatement12",
            "Effect": "Allow",
            "Action": [
                "ecs:RunTask",
                "ecs:StartTask",
                "ecs:StopTask"
            ],
            "Resource": "arn:aws-cn:ecs:*:*:task-definition/*"
        },
        {
            "Sid": "AWSBatchPolicyStatement13",
            "Effect": "Allow",
            "Action": [
                "ecs:StopTask"
            ],
            "Resource": "arn:aws-cn:ecs:*:*:task/*/*"
        },
            "Sid": "AWSBatchPolicyStatement14",
            "Effect": "Allow",
            "Action": [
                "ecs:CreateCluster",
                "ecs:RegisterTaskDefinition"
            ],
            "Resource": "*",
            "Condition": {
                "Null": {
```

```
"aws:RequestTag/AWSBatchServiceTag": "false"
        }
    }
},
{
    "Sid": "AWSBatchPolicyStatement15",
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
        "arn:aws-cn:ec2:*::image/*",
        "arn:aws-cn:ec2:*::snapshot/*",
        "arn:aws-cn:ec2:*:*:subnet/*",
        "arn:aws-cn:ec2:*:*:network-interface/*",
        "arn:aws-cn:ec2:*:*:security-group/*",
        "arn:aws-cn:ec2:*:*:volume/*",
        "arn:aws-cn:ec2:*:*:key-pair/*",
        "arn:aws-cn:ec2:*:*:launch-template/*",
        "arn:aws-cn:ec2:*:*:placement-group/*",
        "arn:aws-cn:ec2:*:*:capacity-reservation/*",
        "arn:aws-cn:ec2:*:*:elastic-gpu/*",
        "arn:aws-cn:elastic-inference:*:*:elastic-inference-accelerator/
        "arn:aws-cn:resource-groups:*:*:group/*"
    1
},
}
    "Sid": "AWSBatchPolicyStatement16",
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": "arn:aws-cn:ec2:*:*:instance/*",
    "Condition": {
        "Null": {
            "aws:RequestTag/AWSBatchServiceTag": "false"
        }
    }
},
    "Sid": "AWSBatchPolicyStatement17",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateTags"
    ],
    "Resource": [
```

```
|
| "Condition": {
| "StringEquals": {
| "ec2:CreateAction": [
| "RunInstances",
| "CreateLaunchTemplate",
| "RequestSpotFleet"
| ]
| }
| }
| }
| }
| }
```

# Amazon managed policy: AWSBatchServiceRole policy

The role permissions policy named **AWSBatchServiceRole** allows Amazon Batch to complete the following actions on specific resources:

The AWSBatchServiceRole managed IAM policy is often used by a role named AWSBatchServiceRole and includes the following permissions. Following the standard security advice of granting least privilege, the AWSBatchServiceRole managed policy can be used as a guide. If any of the permissions that are granted in the managed policy aren't needed for your use case, create a custom policy and add only the permissions that you require. This Amazon Batch managed policy and role can be used with most compute environment types, but service-linked role usage is preferred for a less error prone, better scoped and improved managed experience.

- autoscaling Allows Amazon Batch to create and manage Amazon EC2 Auto Scaling resources. Amazon Batch creates and manages Amazon EC2 Auto Scaling groups for most compute environments.
- ec2 Allows Amazon Batch to manage the lifecycle of Amazon EC2 instances as well as create and manage launch templates and tags. Amazon Batch creates and manages EC2 Spot Fleet requests for some EC2 Spot compute environments.
- ecs Allows Amazon Batch to create and managed Amazon ECS clusters, task definitions and tasks for job execution.
- iam Allows Amazon Batch to validate and pass roles provided by owner to Amazon EC2,
   Amazon EC2 Auto Scaling and Amazon ECS.

• logs – Allows Amazon Batch to create and manage log groups and log streams for Amazon Batch jobs.

**JSON** 

```
}
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AWSBatchPolicyStatement1",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeAccountAttributes",
                "ec2:DescribeInstances",
                "ec2:DescribeInstanceStatus",
                "ec2:DescribeInstanceAttribute",
                "ec2:DescribeSubnets",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeKeyPairs",
                "ec2:DescribeImages",
                "ec2:DescribeImageAttribute",
                "ec2:DescribeSpotInstanceRequests",
                "ec2:DescribeSpotFleetInstances",
                "ec2:DescribeSpotFleetRequests",
                "ec2:DescribeSpotPriceHistory",
                "ec2:DescribeSpotFleetRequestHistory",
                "ec2:DescribeVpcClassicLink",
                "ec2:DescribeLaunchTemplateVersions",
                "ec2:CreateLaunchTemplate",
                "ec2:DeleteLaunchTemplate",
                "ec2:RequestSpotFleet",
                "ec2:CancelSpotFleetRequests",
                "ec2:ModifySpotFleetRequest",
                "ec2:TerminateInstances",
                "ec2:RunInstances",
                "autoscaling:DescribeAccountLimits",
                "autoscaling:DescribeAutoScalingGroups",
                "autoscaling:DescribeLaunchConfigurations",
                "autoscaling:DescribeAutoScalingInstances",
                "autoscaling:DescribeScalingActivities",
                "autoscaling:CreateLaunchConfiguration",
                "autoscaling:CreateAutoScalingGroup",
```

```
"autoscaling:UpdateAutoScalingGroup",
        "autoscaling:SetDesiredCapacity",
        "autoscaling:DeleteLaunchConfiguration",
        "autoscaling:DeleteAutoScalingGroup",
        "autoscaling:CreateOrUpdateTags",
        "autoscaling:SuspendProcesses",
        "autoscaling:PutNotificationConfiguration",
        "autoscaling:TerminateInstanceInAutoScalingGroup",
        "ecs:DescribeClusters",
        "ecs:DescribeContainerInstances",
        "ecs:DescribeTaskDefinition",
        "ecs:DescribeTasks",
        "ecs:ListAccountSettings",
        "ecs:ListClusters",
        "ecs:ListContainerInstances",
        "ecs:ListTaskDefinitionFamilies",
        "ecs:ListTaskDefinitions",
        "ecs:ListTasks",
        "ecs:CreateCluster",
        "ecs:DeleteCluster",
        "ecs:RegisterTaskDefinition",
        "ecs:DeregisterTaskDefinition",
        "ecs:RunTask",
        "ecs:StartTask",
        "ecs:StopTask",
        "ecs:UpdateContainerAgent",
        "ecs:DeregisterContainerInstance",
        "logs:CreateLogGroup",
        "logs:CreateLogStream",
        "logs:PutLogEvents",
        "logs:DescribeLogGroups",
        "iam:GetInstanceProfile",
        "iam:GetRole"
    ],
    "Resource": "*"
},
{
    "Sid": "AWSBatchPolicyStatement2",
    "Effect": "Allow",
    "Action": "ecs:TagResource",
    "Resource": [
        "arn:aws-cn:ecs:*:*:task/*_Batch_*"
    1
},
```

```
{
    "Sid": "AWSBatchPolicyStatement3",
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": [
        11 * 11
    ],
    "Condition": {
        "StringEquals": {
            "iam:PassedToService": [
                 "ec2.amazonaws.com",
                 "ec2.amazonaws.com.cn",
                 "ecs-tasks.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AWSBatchPolicyStatement4",
    "Effect": "Allow",
    "Action": "iam:CreateServiceLinkedRole",
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": [
                 "spot.amazonaws.com",
                 "spotfleet.amazonaws.com",
                 "autoscaling.amazonaws.com",
                 "ecs.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AWSBatchPolicyStatement5",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateTags"
    ],
    "Resource": [
        11 * 11
    ],
    "Condition": {
        "StringEquals": {
```

```
"ec2:CreateAction": "RunInstances"
}
}
}
}
```

# Amazon managed policy: AWSBatchFullAccess

The **AWSBatchFullAccess** policy grants Amazon Batch actions full access to Amazon Batch resources. It also grants describe and list action access for Amazon EC2, Amazon ECS, Amazon EKS, CloudWatch, and IAM services. This is so that IAM identities, either users or roles, can view Amazon Batch managed resources that were created on their behalf. Last, this policy also allows for selected IAM roles to be passed to those services.

You can attach **AWSBatchFullAccess** to your IAM entities. Amazon Batch also attaches this policy to a service role that allows Amazon Batch to perform actions on your behalf.

**JSON** 

```
{
 "Version": "2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Action":[
        "batch: *",
        "cloudwatch:GetMetricStatistics",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeVpcs",
        "ec2:DescribeImages",
        "ec2:DescribeLaunchTemplates",
        "ec2:DescribeLaunchTemplateVersions",
        "ecs:DescribeClusters",
        "ecs:Describe*",
        "ecs:List*",
        "eks:DescribeCluster",
        "eks:ListClusters",
        "logs:Describe*",
```

```
"logs:Get*",
        "logs:TestMetricFilter",
        "logs:FilterLogEvents",
        "iam:ListInstanceProfiles",
        "iam:ListRoles"
      ],
      "Resource":"*"
    },
    {
      "Effect": "Allow",
      "Action":[
        "iam:PassRole"
      ],
      "Resource":[
        "arn:aws-cn:iam::*:role/AWSBatchServiceRole",
        "arn:aws-cn:iam::*:role/service-role/AWSBatchServiceRole",
        "arn:aws-cn:iam::*:role/ecsInstanceRole",
        "arn:aws-cn:iam::*:instance-profile/ecsInstanceRole",
        "arn:aws-cn:iam::*:role/iaws-ec2-spot-fleet-role",
        "arn:aws-cn:iam::*:role/aws-ec2-spot-fleet-role",
        "arn:aws-cn:iam::*:role/AWSBatchJobRole*"
      ]
    },
      "Effect": "Allow",
      "Action":[
        "iam:CreateServiceLinkedRole"
      "Resource": "arn:aws-cn:iam::*:role/*Batch*",
      "Condition": {
        "StringEquals": {
          "iam:AWSServiceName": "batch.amazonaws.com"
          }
      }
    }
 ]
}
```

# Amazon Batch updates to Amazon managed policies

View details about updates to Amazon managed policies for Amazon Batch since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Amazon Batch Document history page.

Change	Description	Date
BatchServiceRolePolicy policy updated	Updated to add support for describing Spot Fleet request history and Amazon EC2 Auto Scaling activities.	December 5, 2023
AWSBatchServiceRole policy added	Updated to add statement IDs, grant Amazon Batch permissions to ec2:DescribeSpotFleetReques tHistory and autoscaling:DescribeScalingActivities.	December 5, 2023
BatchServiceRolePolicy policy updated	Updated to add support for describing Amazon EKS clusters.	October 20, 2022
AWSBatchFullAccess policy updated	Updated to add support for listing and describing Amazon EKS clusters.	October 20, 2022
BatchServiceRolePolicy policy updated	Updated to add support for Amazon EC2 Capacity Reservation groups that are managed by Amazon Resource Groups. For more information, see Work with Capacity Reservation groups in Amazon EC2 User Guide.	May 18, 2022
BatchServiceRolePolicy and AWSBatchS erviceRole policies updated	Updated to add support for describing the status of Amazon Batch managed instances in Amazon EC2 so that unhealthy instances are replaced.	December 6, 2021

Change	Description	Date
BatchServiceRolePolicy policy updated	Updated to add support for placement group, capacity reservation, elastic GPU, and Elastic Inference resources in Amazon EC2.	March 26, 2021
BatchServiceRolePolicy policy added	With the BatchServiceRolePolicy managed policy for the AWSServic eRoleForBatch service-linked role, you can use a service-linked role managed by Amazon Batch. With this policy, you don't need to maintain your own role for use in your compute environments.	March 10, 2021
<u>AWSBatchFullAccess</u> - add permission to add service-linked role	Add IAM permissions to allow the <b>AWSServiceRoleForBatch</b> service-linked role to be added to the account.	March 10, 2021
Amazon Batch started tracking changes	Amazon Batch started tracking changes for its Amazon managed policies.	March 10, 2021

# Amazon Batch IAM policies, roles, and permissions

By default, users don't have permission to create or modify Amazon Batch resources or to perform tasks using the Amazon Batch API, Amazon Batch console, or the Amazon CLI. To allow users to perform these actions, create IAM policies that grant users permission for the specific resources and API operations. Then, attach the policies to the users or groups that require those permissions.

When you attach a policy to a user or group of users, the policy either allows or denies the permissions to perform specific tasks on specific resources. For more information, see <a href="Permissions">Permissions</a> and <a href="Policies">Policies</a> in the IAM User Guide. For more information about managing and creating custom IAM policies, see <a href="Managing IAM Policies">Managing IAM Policies</a>.

Amazon Batch makes calls to other Amazon Web Services services on your behalf. As a result, Amazon Batch must authenticate using your credentials. More specifically, Amazon Batch authenticates by creating an IAM role and policy that provides these permissions. Then, it

associates the role with your compute environments when you create them. For more information, see <u>Amazon ECS instance role</u>, <u>IAM Roles</u>, <u>Using Service-Linked Roles</u>, and <u>Creating a Role to</u> <u>Delegate Permissions to an Amazon Service in the *IAM User Guide*.</u>

### **Topics**

- IAM policy structure
- Resource: Example policies for Amazon Batch
- Resource: Amazon Batch managed policy

# IAM policy structure

The following topics explain the structure of an IAM policy.

#### **Topics**

- Policy syntax
- API actions for Amazon Batch
- Amazon Resource Names for Amazon Batch
- Confirm that users have the required permissions

# **Policy syntax**

An IAM policy is a JSON document that consists of one or more statements. Each statement is structured as follows.

```
{
    "Statement":[{
        "Effect":"effect",
        "Action":"action",
        "Resource":"arn",
        "Condition":{
            "condition":{
            "key":"value"
        }
        }
     }
    }
}
```

IAM policy structure 365

There are four primary elements that make up a statement:

• **Effect**: The *effect* can be Allow or Deny. By default, users don't have permission to use resources and API actions. So, all requests are denied. An explicit allow overrides the default. An explicit deny overrides any allows.

- **Action**: The *action* is the specific API action that you're granting or denying permission for. For instructions on how to specify the *action*, see API actions for Amazon Batch.
- Resource: The resource that's affected by the action. With some Amazon Batch API actions, you can include specific resources in your policy that can be created or modified by the action. To specify a resource in the statement, use its Amazon Resource Name (ARN). For more information, see <a href="Supported resource-level permissions for Amazon Batch API actions">Supported resource-level permissions for Amazon Batch API actions</a> and <a href="Amazon Resource-level permissions">Amazon Resource Names for Amazon Batch</a>. If the Amazon Batch API operation currently doesn't support resource-level permissions, include a wildcard (\*) to specify that all resources can be affected by the action.
- **Condition**: Conditions are optional. They can be used to control when your policy is in effect.

For more information about example IAM policy statements for Amazon Batch, see <u>Resource</u>: Example policies for Amazon Batch.

#### **API actions for Amazon Batch**

In an IAM policy statement, you can specify any API action from any service that supports IAM. For Amazon Batch, use the following prefix with the name of the API action: batch: (for example, batch: SubmitJob and batch: CreateComputeEnvironment).

To specify multiple actions in a single statement, separate each action with a comma.

```
"Action": ["batch:action1", "batch:action2"]
```

You can also specify multiple actions by including a wildcard (\*). For example, you can specify all actions with a name that begins with the word "Describe."

```
"Action": "batch:Describe*"
```

To specify all Amazon Batch API actions, include a wildcard (\*).

```
"Action": "batch:*"
```

IAM policy structure 366

For a list of Amazon Batch actions, see Actions in the Amazon Batch API Reference.

#### Amazon Resource Names for Amazon Batch

Each IAM policy statement applies to the resources that you specify using their Amazon Resource Names (ARNs).

An Amazon Resource Name (ARN) has the following general syntax:

```
arn:aws-cn:[service]:[region]:[account]:resourceType/resourcePath
```

service

The service (for example, batch).

region

The Amazon Web Services Region for the resource (for example, us-east-2). account

The Amazon Web Services account ID, with no hyphens (for example, 123456789012). resourceType

The type of resource (for example, compute-environment).

resourcePath

A path that identifies the resource. You can use a wildcard (\*) in your paths.

Amazon Batch API operations currently support resource-level permissions on several API operations. For more information, see <a href="Supported resource-level permissions for Amazon Batch API actions">Supported resource-level permissions for Amazon Batch API actions</a>. To specify all resources, or if a specific API action doesn't support ARNs, include a wildcard (\*) in the Resource element.

```
"Resource": "*"
```

# Confirm that users have the required permissions

Before you put an IAM policy into production, make sure that it grants users the permissions to use the specific API actions and resources that they need.

IAM policy structure 367

To do this, first create a user for testing purposes and attach the IAM policy to the test user. Then, make a request as the test user. You can make test requests in the console or with the Amazon CLI.



#### Note

You can also test your policies by using the IAM Policy Simulator. For more information about the policy simulator, see Working with the IAM Policy Simulator in the IAM User Guide.

If the policy doesn't grant the user the permissions that you expected, or is overly permissive, you can adjust the policy as needed. Retest until you get the desired results.



#### Important

It can take several minutes for policy changes to propagate before they take effect. Therefore, we recommend that you allow at least five minutes to pass before you test your policy updates.

If an authorization check fails, the request returns an encoded message with diagnostic information. You can decode the message using the DecodeAuthorizationMessage action. For more information, see DecodeAuthorizationMessage in the Amazon Security Token Service API Reference, and decode-authorization-message in the Amazon CLI Command Reference.

# Resource: Example policies for Amazon Batch

You can create specific IAM policies to restrict the calls and resources that users in your account have access to. Then, you can attach those policies to users.

When you attach a policy to a user or group of users, the policy allows or denies the users permission for specific tasks on specific resources. For more information, see Permissions and Policies in the IAM User Guide. For instructions on how to manage and create custom IAM policies, see Managing IAM Policies.

The following examples show policy statements that you can use to control the permissions that users have for Amazon Batch.

### **Examples**

- Resource: Read-only access for Amazon Batch
- Resource: Restrict to POSIX user, Docker image, privilege level, and role on job submission
- Resource: Restrict to job definition prefix on job submission
- Resource: Restrict to a job queue
- Deny action when all conditions match strings
- Resource: Deny action when any condition keys match strings
- Resource: Use the batch:ShareIdentifier condition key

# **Resource: Read-only access for Amazon Batch**

The following policy grants users permissions to use all Amazon Batch API actions with a name that starts with Describe and List.

Unless another statement grants them permission to do so, users don't have permission to perform any actions on the resources. By default, they're denied permission to use API actions.

**JSON** 

# Resource: Restrict to POSIX user, Docker image, privilege level, and role on job submission

The following policy allows a POSIX user to manage their own set of restricted job definitions.

Use the first and second statements to register and deregister any job definition name whose name is prefixed with JobDefA .

The first statement also uses conditional context keys to restrict the POSIX user, privileged status, and container image values within the containerProperties of a job definition. For more information, see RegisterJobDefinition in the Amazon Batch API Reference. In this example, job definitions can only be registered when the POSIX user is set to nobody. The privileged flag is set to false. Last, the image is set to my Image in an Amazon ECR repository.

#### Important

Docker resolves the user parameter to that user uid from within the container image. In most cases, this is found in the /etc/passwd file within the container image. This name resolution can be avoided by using direct uid values in both the job definition and any associated IAM policies. Both the Amazon Batch API operations and the batch: User IAM conditional keys support numeric values.

Use the third statement to restrict to only a specific role to a job definition.

**JSON** 

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "batch:RegisterJobDefinition"
            ],
            "Resource": [
                "arn:aws-cn:batch:<aws_region>:<aws_account_id>:job-definition/
JobDefA_*"
            ],
            "Condition": {
                "StringEquals": {
                     "batch:User": [
                         "nobody"
                     ],
                     "batch:Image": [
```

User Guide Amazon Batch

```
"<aws_account_id>.dkr.ecr.<aws_region>.amazonaws.com/
myImage"
                    ]
                },
                "Bool": {
                     "batch:Privileged": "false"
                }
            }
        },
            "Effect": "Allow",
            "Action": [
                "batch:DeregisterJobDefinition"
            ],
            "Resource": [
                "arn:aws-cn:batch:<aws_region>:<aws_account_id>:job-definition/
JobDefA_*"
            ]
        },
        }
            "Effect": "Allow",
            "Action": [
                "iam:PassRole"
            ],
            "Resource": [
                "arn:aws-cn:iam::<aws_account_id>:role/MyBatchJobRole"
            ]
        }
    ]
}
```

# Resource: Restrict to job definition prefix on job submission

Use the following policy to submit jobs to any job queue with any job definition name that starts with JobDefA.

#### Important

When scoping resource-level access for job submission, you must provide both job queue and job definition resource types.

**JSON** 

```
"Version": "2012-10-17",
  "Statement": [
      "Effect": "Allow",
      "Action": [
        "batch:SubmitJob"
      ],
      "Resource": [
        "arn:aws-cn:batch:<aws_region>:<aws_account_id>:job-definition/
JobDefA_*",
        "arn:aws-cn:batch:<aws_region>:<aws_account_id>:job-queue/*"
    }
  ]
}
```

# Resource: Restrict to a job queue

Use the following policy to submit jobs to a specific job queue that's named queue1 with any job definition name.

#### Important

When scoping resource-level access for job submission, you must provide both job queue and job definition resource types.

**JSON** 

```
"Version": "2012-10-17",
"Statement": [
    "Effect": "Allow",
    "Action": [
      "batch:SubmitJob"
```

```
],
    "Resource": [
        "arn:aws-cn:batch:<aws_region>:<aws_account_id>:job-definition/*",
        "arn:aws-cn:batch:<aws_region>:<aws_account_id>:job-queue/queue1"
        ]
    }
]
```

## Deny action when all conditions match strings

The following policy denies access to the <u>RegisterJobDefinition</u> API operation when both the batch: Image (container image ID) condition key is "string1" and the batch: LogDriver (container log driver) condition key is "string2." Amazon Batch evaluates condition keys on each container. When a job spans multiple containers such as a multi-node parallel job, it's possible for the containers to have different configurations. If multiple condition keys are evaluated in one statement, they're combined using AND logic. So, if any of the multiple condition keys doesn't match for a container, the Deny effect isn't applied for that container. Rather, a different container in the same job might be denied.

For the list of condition keys for Amazon Batch, see <u>Condition keys for Amazon Batch</u> in the <u>Service Authorization Reference</u>. Except for batch: ShareIdentifier, all batch condition keys can be used in this way. The batch: ShareIdentifier condition key is defined for a job, not a job definition.

**JSON** 

```
"Effect": "Deny",
    "Action": "batch:RegisterJobDefinition",
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "batch:Image": "string1",
            "batch:LogDriver": "string2"
        }
    }
}
```

# Resource: Deny action when any condition keys match strings

The following policy denies access to the <u>RegisterJobDefinition</u> API operation when either the batch: Image (container image ID) condition key is "<u>string1</u>" or the batch: LogDriver (container log driver) condition key is "<u>string2</u>." When a job spans multiple containers such as a multi-node parallel job, it's possible for the containers to have different configurations. If multiple condition keys are evaluated in one statement, they're combined using AND logic. So, if any of the multiple condition keys doesn't match for a container, the Deny effect isn't applied for that container. Rather, a different container in the same job might be denied.

For the list of condition keys for Amazon Batch, see <u>Condition keys for Amazon Batch</u> in the <u>Service Authorization Reference</u>. Except for batch: ShareIdentifier, all batch condition keys can be used in this way. (The batch: ShareIdentifier condition key is defined for a job, not a job definition.)

**JSON** 

```
]
    },
      "Effect": "Deny",
      "Action": [
        "batch:RegisterJobDefinition"
      ],
      "Resource": [
        11 * 11
      ],
      "Condition": {
        "StringEquals": {
          "batch:Image": [
             "string1"
          ]
      }
    },
      "Effect": "Deny",
      "Action": [
        "batch:RegisterJobDefinition"
      ],
      "Resource": [
      ],
      "Condition": {
        "StringEquals": {
          "batch:LogDriver": [
            "string2"
          ]
        }
      }
    }
  ]
}
```

# Resource: Use the batch: ShareIdentifier condition key

Use the following policy to submit jobs that use the jobDefA job definition to the jobqueue1 job queue with the lowCpu share identifier.

**JSON** 

```
"Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "batch:SubmitJob"
      ],
      "Resource": [
        "arn:aws-cn::batch:<aws_region>:<aws_account_id>:job-definition/JobDefA",
        "arn:aws-cn::batch:<aws_region>:<aws_account_id>:job-queue/jobqueue1"
      ],
      "Condition": {
        "StringEquals": {
          "batch:ShareIdentifier": [
            "lowCpu"
          1
      }
  ]
}
```

# **Resource: Amazon Batch managed policy**

Amazon Batch provides a managed policy that you can attach to users. This policy provides permission to use Amazon Batch resources and API operations. You can apply this policy directly, or you can use it as a starting point for creating your own policies. For more information about each API operation mentioned in these policies, see Actions in the Amazon Batch API Reference.

#### **AWSBatchFullAccess**

This policy allows full administrator access to Amazon Batch.

**JSON** 

```
{
"Version":"2012-10-17",
```

```
"Statement":[
 {
    "Effect": "Allow",
    "Action":[
      "batch:*",
      "cloudwatch:GetMetricStatistics",
      "ec2:DescribeSubnets",
      "ec2:DescribeSecurityGroups",
      "ec2:DescribeKeyPairs",
      "ec2:DescribeVpcs",
      "ec2:DescribeImages",
      "ec2:DescribeLaunchTemplates",
      "ec2:DescribeLaunchTemplateVersions",
      "ecs:DescribeClusters",
      "ecs:Describe*",
      "ecs:List*",
      "eks:DescribeCluster",
      "eks:ListClusters",
      "logs:Describe*",
      "logs:Get*",
      "logs:TestMetricFilter",
      "logs:FilterLogEvents",
      "iam:ListInstanceProfiles",
     "iam:ListRoles"
   1,
    "Resource":"*"
 },
    "Effect": "Allow",
    "Action":[
      "iam:PassRole"
   ],
    "Resource":[
      "arn:aws-cn:iam::*:role/AWSBatchServiceRole",
      "arn:aws-cn:iam::*:role/service-role/AWSBatchServiceRole",
      "arn:aws-cn:iam::*:role/ecsInstanceRole",
      "arn:aws-cn:iam::*:instance-profile/ecsInstanceRole",
     "arn:aws-cn:iam::*:role/iaws-ec2-spot-fleet-role",
     "arn:aws-cn:iam::*:role/aws-ec2-spot-fleet-role",
      "arn:aws-cn:iam::*:role/AWSBatchJobRole*"
   ]
 },
  {
    "Effect": "Allow",
```

```
"Action":[
    "iam:CreateServiceLinkedRole"
],
    "Resource":"arn:aws-cn:iam::*:role/*Batch*",
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": "batch.amazonaws.com"
        }
    }
}
```

# Amazon Batch IAM execution role

The execution role grants the Amazon ECS container and Amazon Fargate agents permission to make Amazon API calls on your behalf.



The execution role is supported by Amazon ECS container agent version 1.16.0 and later.

The IAM execution role is required depending on the requirements of your task. You can have multiple execution roles for different purposes and services associated with your account.

# Note

For information about the Amazon ECS instance role, see <u>Amazon ECS instance role</u>. For information about service roles, see <u>How Amazon Batch works with IAM</u>.

Amazon ECS provides the AmazonECSTaskExecutionRolePolicy managed policy. This policy contains the required permissions for the common use cases described above. It might be necessary to add inline policies to your execution role for the special use cases outlined below.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```
"Effect": "Allow",
   "Action": [
       "ecr:GetAuthorizationToken",
       "ecr:BatchCheckLayerAvailability",
       "ecr:GetDownloadUrlForLayer",
       "ecr:BatchGetImage",
       "logs:CreateLogStream",
       "logs:PutLogEvents"
    ],
       "Resource": "*"
}
```

# Supported resource-level permissions for Amazon Batch API actions

The term *resource-level permissions* refers to the ability to specify the resources that users are allowed to perform actions on. Amazon Batch has partial support for resource-level permissions. For some Amazon Batch actions, you can control when users are allowed to use those actions based on conditions that must be met. You can also control based on the specific resources that users are allowed to use. For example, you can grant users permissions to submit jobs, but only to a specific job queue and only with a specific job definition.

For details about actions and resource types defined by Amazon Batch, including the format of the ARNs for each of the resource types, see Actions, Resources, and Condition Keys for <u>Amazon Batch</u> in the *Service Authorization Reference*.

# **Tutorial: Create the IAM execution role**

If your account doesn't already have an IAM execution role, use the following steps to create the role.

- 1. Open the IAM console at https://console.amazonaws.cn/iam/.
- 2. In the navigation pane, choose Roles.
- 3. Choose Create role.
- 4. For Trusted entity type, choose Amazon Web Services service.
- 5. For **Service or use case**, choose **Elastic Container Service**. Then choose **Elastic Container Service Task** again.
- 6. Choose **Next**.

- 7. For **Permissions policies**, search for **AmazonECSTaskExecutionRolePolicy**.
- 8. Choose the check box to the left of the **AmazonECSTaskExecutionRolePolicy** policy, and then choose **Next**.

9. For **Role Name**, enter ecsTaskExecutionRole and then choose **Create role**.

### Tutorial: Check the IAM execution role

Use the following procedure to check that your account already has the IAM execution role and attach the managed IAM policy, if needed.

- 1. Open the IAM console at https://console.amazonaws.cn/iam/.
- 2. In the navigation pane, choose **Roles**.
- 3. Search the list of roles for ecsTaskExecutionRole. If you can't find the role, see <u>Tutorial</u>: <u>Create the IAM execution role</u>. If you found the role, choose the role to view the attached policies.
- 4. On the **Permissions** tab, verify that the **AmazonECSTaskExecutionRolePolicy** managed policy is attached to the role. If the policy is attached, your execution role is properly configured. If not, follow the substeps below to attach the policy.
  - a. Choose Add permissions, then choose Attach policies.
  - b. Search for AmazonECSTaskExecutionRolePolicy.
  - c. Check the box to the left of the **AmazonECSTaskExecutionRolePolicy** policy and choose **Attach policies**.
- 5. Choose Trust relationships.
- 6. Verify that the trust relationship contains the following policy. If the trust relationship matches the policy below, the role is configured correctly. If the trust relationship does not match, choose **Edit trust policy**, enter the following, and choose **Update policy**.

**JSON** 

```
{
  "Version": "2012-10-17",
  "Statement": [
     {
        "Sid": "",
```

```
"Effect": "Allow",
    "Principal": {
        "Service": "ecs-tasks.amazonaws.com"
     },
      "Action": "sts:AssumeRole"
     }
]
```

#### Use service-linked roles for Amazon Batch

Amazon Batch uses Amazon Identity and Access Management (IAM) <u>service-linked roles</u>. A service-linked role is a unique type of IAM role that is linked directly to Amazon Batch. Service-linked roles are predefined by Amazon Batch and include all the permissions that the service requires to call other Amazon services on your behalf.

A service-linked role makes setting up Amazon Batch easier because you don't have to manually add the necessary permissions. Amazon Batch defines the permissions of its service-linked roles, and unless defined otherwise, only Amazon Batch can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

### Note

Do one of the following to specify a service role for an Amazon Batch compute environment.

- Use an empty string for the service role. This lets Amazon Batch create the service role.
- Specify the service role in the following format: arn:awscn:iam::account\_number:role/aws-service-role/batch.amazonaws.com/ AWSServiceRoleForBatch.

For more information, see <u>Incorrect role name or ARN</u> in the Amazon Batch User Guide.

You can delete a service-linked role only after first deleting their related resources. This protects your Amazon Batch resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see <u>Amazon Services That</u> <u>Work with IAM</u> and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

# Service-linked role permissions for Amazon Batch

Amazon Batch uses the service-linked role named **AWSServiceRoleForBatch**. The **AWSServiceRoleForBatch** role allows Amazon Batch to create and manage Amazon resources on your behalf.

The **AWSServiceRoleForBatch** service-linked role trusts the batch.amazonaws.com service principal to assume the role.

The IAM policy named <u>BatchServiceRolePolicy</u> allows Amazon Batch to complete the following actions on specific resources:

- autoscaling Allows Amazon Batch to create and manage Amazon EC2 Auto Scaling resources. Amazon Batch creates and manages Amazon EC2 Auto Scaling groups for most compute environments.
- ec2 Allows Amazon Batch to control the lifecycle of Amazon EC2 instances as well as create
  and manage launch templates and tags. Amazon Batch creates and manages EC2 Spot Fleet
  requests for some EC2 Spot compute environments.
- ecs Allows Amazon Batch to create and managed Amazon ECS clusters, task-definitions and tasks for job execution.
- eks Allows Amazon Batch to describe the Amazon EKS cluster resource for validations.
- iam Allows Amazon Batch to validate and pass roles provided by owner to Amazon EC2,
   Amazon EC2 Auto Scaling and Amazon ECS.
- logs Allows Amazon Batch to create and manage log groups and log streams for Amazon Batch jobs.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see <a href="Service-Linked Role Permissions">Service-Linked Role Permissions</a> in the IAM User Guide.

# Creating a service-linked role for Amazon Batch

You don't need to manually create a service-linked role. When you CreateComputeEnvironment in the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API, and don't

specify a value for the serviceRole parameter, Amazon Batch creates the service-linked role for you.

#### Important

This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. Also, if you were using the Amazon Batch service before March 10, 2021, when it began supporting service-linked roles, then Amazon Batch created the AWSServiceRoleForBatch role in your account. To learn more, see A New Role Appeared in My IAM Account.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you CreateComputeEnvironment, Amazon Batch creates the service-linked role for you again.

# **Editing a service-linked role for Amazon Batch**

With Amazon Batch, you can't edit the AWSServiceRoleForBatch service-linked role. After you create a service-linked role, you can't change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Editing a Service-Linked Role in the IAM User Guide.

# To allow an IAM entity to edit the description of the AWSServiceRoleForBatch service-linked role

Add the following statement to the permissions policy. This allows the IAM entity to edit the description of a service-linked role.

```
{
    "Effect": "Allow",
    "Action": [
        "iam:UpdateRoleDescription"
    "Resource": "arn:aws-cn:iam::*:role/aws-service-role/batch.amazonaws.com/
AWSServiceRoleForBatch",
    "Condition": {"StringLike": {"iam:AWSServiceName": "batch.amazonaws.com"}}
}
```

# Deleting a service-linked role for Amazon Batch

We recommend, if you no longer need to use a feature or service that requires a service-linked role, you delete that role. That way, you don't have an unused entity that's not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

#### To allow an IAM entity to delete the AWSServiceRoleForBatch service-linked role

Add the following statement to the permissions policy. This allows the IAM entity to delete a service-linked role.

```
{
    "Effect": "Allow",
    "Action": Γ
        "iam:DeleteServiceLinkedRole",
        "iam:GetServiceLinkedRoleDeletionStatus"
    ],
    "Resource": "arn:aws-cn:iam::*:role/aws-service-role/batch.amazonaws.com/
AWSServiceRoleForBatch",
    "Condition": {"StringLike": {"iam:AWSServiceName": "batch.amazonaws.com"}}
}
```

# Cleaning up a service-linked role

Before you can use IAM to delete a service-linked role, you must first confirm that the role has no active sessions and delete all of the Amazon Batch compute environments that use the role in all Amazon Regions in a single partition.

#### To check whether the service-linked role has an active session

- 1. Open the IAM console at https://console.amazonaws.cn/iam/.
- 2. In the navigation pane, choose Roles and then the AWSServiceRoleForBatch name (not the check box).
- On the **Summary** page, choose **Access Advisor** and review recent activity for the service-linked role.



#### Note

If you don't know whether Amazon Batch is using the AWSServiceRoleForBatch role, you can try to delete the role. If the service is using the role, then the role will fail to

delete. You can view the Regions where the role is being used. If the role is being used, then you must wait for the session to end before you can delete the role. You can't revoke the session for a service-linked role.

#### To remove Amazon Batch resources used by the AWSServiceRoleForBatch service-linked role

You must delete all Amazon Batch compute environments that use the AWSServiceRoleForBatch role in all Amazon Regions before you can delete the AWSServiceRoleForBatch role.

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Region to use.
- 3. In the navigation pane, choose **Compute environments**.
- 4. Select the compute environment.
- 5. Choose **Disable**. Wait for the **State** to change to **DISABLED**.
- 6. Select the compute environment.
- 7. Choose **Delete**. Confirm that you want to delete the compute environment by choosing **Delete** compute environment.
- 8. Repeat steps 1–7 for all compute environments that use the service-linked role in all Regions.

## Deleting a service-linked role in IAM (Console)

You can use the IAM console to delete a service-linked role.

#### To delete a service-linked role (console)

- 1. Sign in to the Amazon Web Services Management Console and open the IAM console at https://console.amazonaws.cn/iam/.
- 2. In the navigation pane of the IAM console, choose **Roles**. Then select the check box next to AWSServiceRoleForBatch, not the name or row itself.
- 3. Choose **Delete role**.
- 4. In the confirmation dialog box, review the service last accessed data, which shows when each of the selected roles last accessed an Amazon Web Services service. This helps you to confirm whether the role is currently active. If you want to proceed, choose **Yes, Delete** to submit the service-linked role for deletion.

Use service-linked roles 385

Watch the IAM console notifications to monitor the progress of the service-linked role deletion. Because the IAM service-linked role deletion is asynchronous, after you submit the role for deletion, the deletion task can succeed or fail.

- If the task succeeds, then the role is removed from the list and a notification of success appears at the top of the page.
- If the task fails, you can choose **View details** or **View Resources** from the notifications to learn why the deletion failed. If the deletion fails because the role is using the service's resources, then the notification includes a list of resources, if the service returns that information. You can then clean up the resources and submit the deletion again.

#### Note

You might have to repeat this process several times, depending on the information that the service returns. For example, your service-linked role might use six resources and your service might return information about five of them. If you clean up the five resources and submit the role for deletion again, the deletion fails and the service reports the one remaining resource. A service might return all of the resources, a few of them, or it might not report any resources.

 If the task fails and the notification does not include a list of resources, then the service might not return that information. To learn how to clean up the resources for that service, see Amazon Web Services services that work with IAM. Find your service in the table, and choose the **Yes** link to view the service-linked role documentation for that service.

## Deleting a service-linked role in IAM (Amazon CLI)

You can use IAM commands from the Amazon Command Line Interface to delete a service-linked role.

#### To delete a service-linked role (CLI)

Because a service-linked role can't be deleted if it's being used or has associated resources, you must submit a deletion request. That request can be denied if these conditions aren't met. You must capture the deletion-task-id from the response to check the status of the deletion task. Enter the following command to submit a service-linked role deletion request:

\$ aws iam delete-service-linked-role --role-name AWSServiceRoleForBatch

Use service-linked roles 386

Use the following command to check the status of the deletion task: 2.

\$ aws iam get-service-linked-role-deletion-status --deletion-task-id deletion-taskid

The status of the deletion task can be NOT\_STARTED, IN\_PROGRESS, SUCCEEDED, or FAILED. If the deletion fails, the call returns the reason that it failed so that you can troubleshoot. If the deletion fails because the role is using the service's resources, then the notification includes a list of resources, if the service returns that information. You can then clean up the resources and submit the deletion again.



### Note

You might have to repeat this process several times, depending on the information that the service returns. For example, your service-linked role might use six resources and your service might return information about five of them. If you clean up the five resources and submit the role for deletion again, the deletion fails and the service reports the one remaining resource. A service might return all of the resources, a few of them. Or, it might not report any resources. To learn how to clean up the resources for a service that doesn't report any resources, see Amazon services that work with IAM. Find your service in the table, and choose the Yes link to view the service-linked role documentation for that service.

## Deleting a service-linked role in IAM (AmazonAPI)

You can use the IAM API to delete a service-linked role.

#### To delete a service-linked role (API)

To submit a deletion request for a service-linked roll, call DeleteServiceLinkedRole. In the 1. request, specify the AWSServiceRoleForBatch role name.

Because a service-linked role cannot be deleted if it is being used or has associated resources, you must submit a deletion request. That request can be denied if these conditions are not met. You must capture the DeletionTaskId from the response to check the status of the deletion task.

Use service-linked roles 387

To check the status of the deletion, call GetServiceLinkedRoleDeletionStatus. In the request, 2. specify the DeletionTaskId.

The status of the deletion task can be NOT\_STARTED, IN\_PROGRESS, SUCCEEDED, or FAILED. If the deletion fails, the call returns the reason that it failed so that you can troubleshoot. If the deletion fails because the role is using the service's resources, then the notification includes a list of resources, if the service returns that information. You can then clean up the resources and submit the deletion again.



### Note

You might have to repeat this process several times, depending on the information that the service returns. For example, your service-linked role might use six resources and your service might return information about five of them. If you clean up the five resources and submit the role for deletion again, the deletion fails and the service reports the one remaining resource. A service might return all of the resources, a few of them, or it might not report any resources. To learn how to clean up the resources for a service that does not report any resources, see Amazon Web Services services that work with IAM. Find your service in the table, and choose the Yes link to view the service-linked role documentation for that service.

## Supported Regions for Amazon Batch service-linked roles

Amazon Batch supports using service-linked roles in all of the Regions where the service is available. For more information, see Amazon Batch endpoints.

## Amazon ECS instance role

Amazon Batch compute environments are populated with Amazon ECS container instances. They run the Amazon ECS container agent locally. The Amazon ECS container agent makes calls to various Amazon API operations on your behalf. Therefore, container instances that run the agent require an IAM policy and role for these services to recognize that the agent belongs to you. You must create an IAM role and an instance profile for the container instances to use when they're launched. Otherwise, you can't create a compute environment and launch container instances into it. This requirement applies to container instances launched with or without the Amazon ECS optimized AMI provided by Amazon. For more information, see Amazon ECS instance role in the Amazon Elastic Container Service Developer Guide.

#### **Topics**

Tutorial: Check your account's Amazon ECS instance role

## **Tutorial: Check your account's Amazon ECS instance role**

The Amazon ECS instance role and instance profile are automatically created for you in the console first-run experience. However, you can follow these steps to check if your account already has the Amazon ECS instance role and instance profile. The following steps also cover how to attach the managed IAM policy.

#### Tutorial: Check for the ecsInstanceRole in the IAM console

- Open the IAM console at https://console.amazonaws.cn/iam/. 1.
- In the navigation pane, choose **Roles**. 2.
- 3. Search the list of roles for ecsInstanceRole. If the role doesn't exist, use the following steps to create the role.
  - Choose Create Role. a.
  - For Trusted entity type, choose Amazon Web Services service. b.
  - For **Common use cases**, choose **EC2**. C.
  - d. Choose **Next**.
  - For **Permissions policies**, search for **AmazonEC2ContainerServiceforEC2Role**. e.
  - f. Choose the check box next to **AmazonEC2ContainerServiceforEC2Role**, then choose Next.
  - For **Role Name**, type ecsInstanceRole and choose **Create Role**.



#### Note

If you use the Amazon Web Services Management Console to create a role for Amazon EC2, the console creates an instance profile with the same name as the role.

Alternatively, you can use the Amazon CLI to create the ecsInstanceRole IAM role. The following example creates an IAM role with a trust policy and an Amazon managed policy.

#### Tutorial: Create an IAM role and instance profile (Amazon CLI)

 Create the following trust policy and save it in a text file that's named ecsInstanceRolerole-trust-policy.json.

**JSON** 

2. Use the <u>create-role</u> command to create the ecsInstanceRole role. Specify the trust policy file location in the assume-role-policy-document parameter.

```
$ aws iam create-role \
    --role-name ecsInstanceRole \
    --assume-role-policy-document file://ecsInstanceRole-role-trust-policy.json
```

The following is an example response.

```
"Action": "sts:AssumeRole",
         }
      ]
 }
}
```

3. Use the create-instance-profile command to create an instance profile that's named ecsInstanceRole.



#### Note

You need to create roles and instance profiles as separate actions in the Amazon CLI and Amazon API.

```
$ aws iam create-instance-profile --instance-profile-name ecsInstanceRole
```

The following is an example response.

```
{
    "InstanceProfile": {
        "Path": "/",
        "InstanceProfileName": "ecsInstanceRole",
        "InstanceProfileId": "AIPAT46P5RDITREXAMPLE",
        "Arn": "arn:aws-cn:iam::123456789012:instance-profile/ecsInstanceRole",
        "CreateDate": "2022-06-30T23:53:34.093Z",
        "Roles": [],
}
```

4. Use the add-role-to-instance-profile command to add the ecsInstanceRole role to the ecsInstanceRole instance profile.

```
aws iam add-role-to-instance-profile \
    --role-name ecsInstanceRole --instance-profile-name ecsInstanceRole
```

5. Use the attach-role-policy command to attach the AmazonEC2ContainerServiceforEC2Role Amazon managed policy to the ecsInstanceRole role.

```
$ aws iam attach-role-policy \
    --policy-arn arn:aws-cn:iam::aws:policy/service-role/
AmazonEC2ContainerServiceforEC2Role \
    --role-name ecsInstanceRole
```

## Amazon EC2 spot fleet role

If you create a managed compute environment that uses Amazon EC2 Spot Fleet Instances, you must create the AmazonEC2SpotFleetTaggingRole policy. This policy grants Spot Fleet permission to launch, tag, and terminate instances on your behalf. Specify the role in your Spot Fleet request. You must also have the AWSServiceRoleForEC2Spot and AWSServiceRoleForEC2SpotFleet service-linked roles for Amazon EC2 Spot and Spot Fleet. Use the following instruction to create all of these roles. For more information, see <a href="Using Service-Linked Roles">Using Service-Linked Roles</a> and <a href="Creating a Role to Delegate Permissions to an Amazon Service">Creating a Role to Delegate Permissions to an Amazon Service</a> in the IAM User Guide.

### **Topics**

- Tutorial: Create Amazon EC2 spot fleet roles in the Amazon Web Services Management Console
- Tutorial: Create Amazon EC2 spot fleet roles with the Amazon CLI

## Tutorial: Create Amazon EC2 spot fleet roles in the Amazon Web Services Management Console

To create the AmazonEC2SpotFleetTaggingRole IAM service-linked role for Amazon EC2 Spot Fleet

- 1. Open the IAM console at <a href="https://console.amazonaws.cn/iam/">https://console.amazonaws.cn/iam/</a>.
- 2. For Access Management, choose Roles,
- 3. For Roles, choose Create role.
- 4. From Select trusted entity for Trusted entity type, choose Amazon Web Services service.
- For Use cases for other Amazon Web Services services, choose EC2 and then choose EC2 Spot Fleet Tagging.
- Choose Next.
- 7. From **Permissions policies** for **Policy name**, verify AmazonEC2SpotFleetTaggingRole.

Amazon EC2 spot fleet role 392

- Choose Next.
- 9. For Name, review, and create:
  - a. For **Role name**, enter a name to identify the role.
  - b. For **Description**, enter a short explanation for the policy.
  - c. (Optional) For **Step 1: Select trusted entities**, choose **Edit** to modify the code.
  - d. (Optional) For **Step 2: Add permissions**, choose **Edit** to modify the code.
  - e. (Optional) For **Add tags**, choose **Add tag** to add tags to the resource.
  - f. Choose Create role.

#### Note

In the past, there were two managed policies for the Amazon EC2 Spot Fleet role.

- AmazonEC2SpotFleetRole: This is the original managed policy for the Spot Fleet role. However, we no longer recommend that you use it with Amazon Batch. This policy doesn't support Spot Instance tagging in compute environments, which is required to use the AWSServiceRoleForBatch service-linked role. If you previously created a Spot Fleet role with this policy, apply the new recommended policy to that role. For more information, see Spot Instances not tagged on creation.
- AmazonEC2SpotFleetTaggingRole: This role provides all of the necessary permissions to tag Amazon EC2 Spot Instances. Use this role to allow Spot Instance tagging on your Amazon Batch compute environments.

## Tutorial: Create Amazon EC2 spot fleet roles with the Amazon CLI

To create the AmazonEC2SpotFleetTaggingRole IAM role for your Spot Fleet compute environments

1. Run the following command with the Amazon CLI.

```
$ aws iam create-role --role-name AmazonEC2SpotFleetTaggingRole \
    --assume-role-policy-document '{
    "Version":"2012-10-17",
    "Statement":[
        {
```

Amazon EC2 spot fleet role 393

```
"Sid":"",
    "Effect":"Allow",
    "Principal": {
        "Service":"spotfleet.amazonaws.com"
        },
        "Action":"sts:AssumeRole"
     }
]
]
```

2. To attach the **AmazonEC2SpotFleetTaggingRole** managed IAM policy to your **AmazonEC2SpotFleetTaggingRole** role, run the following command with the Amazon CLI.

```
$ aws iam attach-role-policy \
    --policy-arn \
    arn:aws-cn:iam::aws:policy/service-role/AmazonEC2SpotFleetTaggingRole \
    --role-name \
    AmazonEC2SpotFleetTaggingRole
```

#### To create the AWSServiceRoleForEC2Spot IAM service-linked role for Amazon EC2 Spot

## Note

If the AWSServiceRoleForEC2Spot IAM service-linked role already exists, you see an error message that resembles the following.

```
An error occurred (InvalidInput) when calling the CreateServiceLinkedRole operation:

Service role name AWSServiceRoleForEC2Spot has been taken in this account, please try a different suffix.
```

Run the following command with the Amazon CLI.

```
$ aws iam create-service-linked-role --aws-service-name spot.amazonaws.com
```

Amazon EC2 spot fleet role 394

## To create the AWSServiceRoleForEC2SpotFleet IAM service-linked role for Amazon EC2 **Spot Fleet**



#### Note

If the AWSServiceRoleForEC2SpotFleet IAM service-linked role already exists, you see an error message that resembles the following.

An error occurred (InvalidInput) when calling the CreateServiceLinkedRole

Service role name AWSServiceRoleForEC2SpotFleet has been taken in this account, please try a different suffix.

Run the following command with the Amazon CLI.

aws iam create-service-linked-role --aws-service-name spotfleet.amazonaws.com

## **EventBridge IAM role**

Amazon EventBridge delivers a near-real time stream of system events that describe changes in Amazon resources. Amazon Batch jobs are available as EventBridge targets. Using simple rules that you can quickly set up, you can match events and submit Amazon Batch jobs in response to them. Before you can submit Amazon Batch jobs with EventBridge rules and targets, EventBridge must have permissions to run Amazon Batch jobs on your behalf.



#### Note

When you create a rule in the EventBridge console that specifies an Amazon Batch queue as a target, you can create this role. For an example walkthrough, see Amazon Batch jobs as EventBridge targets. You can create the EventBridge role manually using the IAM console. For instructions, see Creating a role using custom trust policies (console) in the IAM User Guide.

The trust relationship for your EventBridge IAM role must provide the events.amazonaws.com service principal the ability to assume the role.

EventBridge IAM role 395

**JSON** 

Make sure that the policy that's attached to your EventBridge IAM role allows batch: SubmitJob permissions on your resources. In the following example, Amazon Batch provides the AWSBatchServiceEventTargetRole managed policy to provide these permissions.

**JSON** 

EventBridge IAM role 396

## Create a virtual private cloud

Compute resources in your compute environments need external network access to communicate with Amazon Batch and Amazon ECS service endpoints. However, you might have jobs that you want to run in private subnets. To have the flexibility to run jobs in either a public or private subnet, create a VPC that has both public and private subnets.

You can use Amazon Virtual Private Cloud (Amazon VPC) to launch Amazon resources into a virtual network that you define. This topic provides a link to the Amazon VPC wizard and a list of the options to select.

#### Create a VPC

For information about how to create an Amazon VPC, see <u>Create a VPC only</u> in the *Amazon VPC User Guide* and use the following table to determine what options to select.

Option	Value
Resources to create	VPC only
Name	Optionally provide a name for your VPC.
IPv4 CIDR block	IPv4 CIDR manual input  The CIDR block size must have a size between /16 and /28.
IPv6 CIDR block	No IPv6 CIDR block
Tenancy	Default

For more information about Amazon VPC, see What is Amazon VPC? in the Amazon VPC User Guide.

## **Next steps**

After you have created your VPC, consider the following next steps:

Create a virtual private cloud 397

• Create security groups for your public and private resources if they require inbound network access. For more information, see Work with security groups in the *Amazon VPC User Guide*.

- Create an Amazon Batch managed compute environment that launches compute resources into your new VPC. For more information, see <a href="Create a compute environment">Create a compute environment</a>. If you use the compute environment creation wizard in the Amazon Batch console, you can specify the VPC that you just created and the public or private subnets that you want to launch your instances into.
- Create an Amazon Batch job queue that's mapped to your new compute environment. For more information, see Create a job queue.
- Create a job definition to run your jobs with. For more information, see <u>Create a single-node job</u> definition .
- Submit a job with your job definition to your new job queue. This job lands in the compute environment that you created with your new VPC and subnets. For more information, see Tutorial: submit a job.

## Use an interface endpoint to Access Amazon Batch

You can use Amazon PrivateLink to create a private connection between your VPC and Amazon Batch. You can access Amazon Batch as if it were in your VPC, without the use of an internet gateway, NAT device, VPN connection, or Amazon Direct Connect connection. Instances in your VPC don't need public IP addresses to access Amazon Batch.

You establish this private connection by creating an *interface endpoint*, powered by Amazon PrivateLink. We create an endpoint network interface in each subnet that you enable for the interface endpoint. These are requester-managed network interfaces that serve as the entry point for traffic destined for Amazon Batch.

For more information, see Interface VPC endpoints in the Amazon PrivateLink Guide.

## **Considerations for Amazon Batch**

Before you set up an interface endpoint for Amazon Batch, review <u>Interface endpoint properties</u> and <u>limitations</u> in the *Amazon PrivateLink Guide*.

Amazon Batch supports making calls to all of its API actions through the interface endpoint.

Before you set up interface VPC endpoints for Amazon Batch, be aware of the following considerations:

VPC endpoints 398

 Jobs using Fargate resources launch type don't require the interface VPC endpoints for Amazon ECS, but you might need interface VPC endpoints for Amazon Batch, Amazon ECR, Secrets Manager, or Amazon CloudWatch Logs described in the following points.

- To run jobs, you must create the interface VPC endpoints for Amazon ECS. For more information, see <a href="Interface VPC Endpoints">Interface VPC Endpoints (Amazon PrivateLink)</a> in the Amazon Elastic Container Service Developer Guide.
- To allow your jobs to pull private images from Amazon ECR, you must create the interface VPC endpoints for Amazon ECR. For more information, see <u>Interface VPC Endpoints (Amazon PrivateLink)</u> in the *Amazon Elastic Container Registry User Guide*.
- To allow your jobs to pull sensitive data from Secrets Manager, you must create the interface VPC endpoints for Secrets Manager. For more information, see <u>Using Secrets Manager with</u> VPC Endpoints in the *Amazon Secrets Manager User Guide*.
- If your VPC doesn't have an internet gateway and your jobs use the awslogs log driver to send log information to CloudWatch Logs, you must create an interface VPC endpoint for CloudWatch Logs. For more information, see <u>Using CloudWatch Logs with Interface VPC</u> <u>Endpoints</u> in the *Amazon CloudWatch Logs User Guide*.
- Jobs using the EC2 resources require that the container instances that they're launched on to run version 1.25.1 or later of the Amazon ECS container agent. For more information, see <u>Amazon ECS Linux container agent versions</u> in the *Amazon Elastic Container Service Developer Guide*.
- VPC endpoints currently don't support cross-Region requests. Ensure that you create your endpoint in the same Region where you plan to issue your API calls to Amazon Batch.
- VPC endpoints only support Amazon-provided DNS through Amazon Route 53. If you want to
  use your own DNS, you can use conditional DNS forwarding. For more information, see <a href="DHCP">DHCP</a>
  Options Sets in the Amazon VPC User Guide.
- The security group attached to the VPC endpoint must allow incoming connections on port 443 from the private subnet of the VPC.
- Amazon Batch does not support VPC interface endpoints in the following Amazon Web Services Regions:
  - Asia Pacific (Osaka) (ap-northeast-3)
  - Asia Pacific (Jakarta) (ap-southeast-3)

Considerations 399

## Create an interface endpoint for Amazon Batch

You can create an interface endpoint for Amazon Batch using either the Amazon VPC console or the Amazon Command Line Interface (Amazon CLI). For more information, see <u>Create an interface endpoint</u> in the *Amazon PrivateLink Guide*.

Create an interface endpoint for Amazon Batch using the following service name:

```
com.amazonaws.region.batch
```

#### For example:

```
com.amazonaws.us-east-2.batch
```

In the aws-cn partition, the format is different:

```
cn.com.amazonaws.region.batch
```

#### For example:

```
cn.com.amazonaws.cn-northwest-1.batch
```

If you enable private DNS for the interface endpoint, you can make API requests to Amazon Batch using its default Regional DNS name. For example, batch.us-east-1.amazonaws.com.

For more information, see <u>Access a service through an interface endpoint</u> in the *Amazon PrivateLink Guide*.

## Create an endpoint policy for your interface endpoint

An endpoint policy is an IAM resource that you can attach to an interface endpoint. The default endpoint policy allows full access to Amazon Batch through the interface endpoint. To control the access allowed to Amazon Batch from your VPC, attach a custom endpoint policy to the interface endpoint.

An endpoint policy specifies the following information:

- The principals that can perform actions (Amazon Web Services accounts, users, and IAM roles).
- The actions that can be performed.
- The resources on which the actions can be performed.

Create an interface endpoint 400

For more information, see <u>Control access to services using endpoint policies</u> in the *Amazon PrivateLink Guide*.

#### **Example: VPC endpoint policy for Amazon Batch actions**

The following is an example of a custom endpoint policy. When you attach this policy to your interface endpoint, it grants access to the listed Amazon Batch actions for all principals on all resources.

## **Compliance validation for Amazon Batch**

To learn whether an Amazon Web Services service is within the scope of specific compliance programs, see <u>Amazon Web Services services in Scope by Compliance Program</u> and choose the compliance program that you are interested in. For general information, see <u>Amazon Web Services</u> Compliance Programs.

You can download third-party audit reports using Amazon Artifact. For more information, see Downloading Reports in Amazon Artifact.

Your compliance responsibility when using Amazon Web Services services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

• <u>Security & Compliance</u> – These solution implementation guides discuss architectural considerations and provide steps for deploying security and compliance features.

Compliance validation 401

 <u>Amazon Compliance Resources</u> – This collection of workbooks and guides might apply to your industry and location.

- <u>Evaluating Resources with Rules</u> in the *Amazon Config Developer Guide* The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>Amazon Security Hub</u> This Amazon Web Services service provides a comprehensive view of
  your security state within Amazon. Security Hub uses security controls to evaluate your Amazon
  resources and to check your compliance against security industry standards and best practices.
   For a list of supported services and controls, see Security Hub controls reference.
- Amazon GuardDuty This Amazon Web Services service detects potential threats to your
  Amazon Web Services accounts, workloads, containers, and data by monitoring your
  environment for suspicious and malicious activities. GuardDuty can help you address various
  compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated
  by certain compliance frameworks.

## Infrastructure security in Amazon Batch

As a managed service, Amazon Batch is protected by Amazon global network security. For information about Amazon security services and how Amazon protects infrastructure, see <a href="Amazon Cloud Security">Amazon Security</a>. To design your Amazon environment using the best practices for infrastructure security, see <a href="Infrastructure Protection">Infrastructure Protection</a> in Security Pillar Amazon Well-Architected Framework.

You use Amazon published API calls to access Amazon Batch through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the <u>Amazon Security Token Service</u> (Amazon STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Amazon Batch does support resource-based access policies, which can include restrictions based on the source IP address.

Infrastructure security 402

You can also use Amazon Batch policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given Amazon Batch resource from only the specific VPC within the Amazon network.

## **Cross-service confused deputy prevention**

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more privileged entity to perform the action. In Amazon, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, Amazon provides tools that help you protect your data for all services with service principals that have been given access to resources in your account.

We recommend using the <a href="mailto:aws:SourceArn">aws:SourceAccount</a> global condition context keys in resource policies to limit the permissions that Amazon Batch gives another service to the resource. If the <a href="mailto:aws:SourceArn">aws:SourceArn</a> value does not contain the account ID, such as an Amazon S3 bucket ARN, you must use both global condition context keys to limit permissions. If you use both global condition context keys and the <a href="mailto:aws:SourceArn">aws:SourceArn</a> value contains the account ID, the <a href="mailto:aws:SourceArn">aws:SourceArn</a> value must use the same account ID when used in the same policy statement. Use <a href="mailto:aws:SourceArn">aws:SourceArn</a> if you want only one resource to be associated with the cross-service access. Use <a href="mailto:aws:SourceAccount">aws:SourceAccount</a> if you want to allow any resource in that account to be associated with the cross-service use.

The value of aws: SourceArn must be the resource that Amazon Batch stores.

The most effective way to protect against the confused deputy problem is to use the aws:SourceArn global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the aws:SourceArn global context condition key with wildcard characters (\*) for the unknown portions of the ARN. For example, arn:aws-cn:servicename:\*:123456789012:\*.

The following examples shows how you can use the aws:SourceArn and aws:SourceAccount global condition context keys in Amazon Batch to prevent the confused deputy problem.

## Example: Role for accessing only one compute environment

The following role can only be used to access one compute environment. The job name must be specified as \* because the job queue can be associated with multiple compute environments.

**JSON** 

```
"Version": "2012-10-17",
  "Statement": [
      "Effect": "Allow",
      "Principal": {
        "Service": "batch.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "123456789012"
        },
        "ArnLike": {
          "aws:SourceArn": [
            "arn:aws-cn:batch:us-east-1:123456789012:compute-environment/testCE",
            "arn:aws-cn:batch:us-east-1:123456789012:job/*"
          ]
        }
      }
    }
  ]
}
```

## **Example: Role for accessing multiple compute environments**

The following role can be used to access multiple compute environments. The job name must be specified as \* because the job queue can be associated with multiple compute environments.

**JSON** 

```
{
```

```
"Version": "2012-10-17",
  "Statement": [
      "Effect": "Allow",
      "Principal": {
        "Service": "batch.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "123456789012"
        },
        "ArnLike": {
          "aws:SourceArn": [
            "arn:aws-cn:batch:us-east-1:123456789012:compute-environment/*",
            "arn:aws-cn:batch:us-east-1:123456789012:job/*"
          1
        }
      }
    }
  ]
}
```

## Logging Amazon Batch API calls with Amazon CloudTrail

Amazon Batch is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, role, or an Amazon service in Amazon Batch. CloudTrail captures all API calls for Amazon Batch as events. The calls captured include calls from the Amazon Batch console and code calls to the Amazon Batch API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Batch. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Batch, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the Amazon CloudTrail User Guide.

#### **Topics**

- Amazon Batch information in CloudTrail
- Reference: Understanding Amazon Batch log file entries

CloudTrail 405

## Amazon Batch information in CloudTrail

CloudTrail is enabled on your Amazon account when you create the account. When activity occurs in Amazon Batch, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search, and download recent events in your Amazon account. For more information, see Viewing Events with CloudTrail Event History.

For an ongoing record of events in your Amazon account, including events for Amazon Batch, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all Amazon Regions. The trail logs events from all Regions in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- Overview for Creating a Trail
- CloudTrail Supported Services and Integrations
- Configuring Amazon SNS Notifications for CloudTrail
- Receiving CloudTrail Log Files from Multiple Regions and Receiving CloudTrail Log Files from Multiple Accounts

All Amazon Batch actions are logged by CloudTrail and are documented in the https://docs.amazonaws.cn/batch/latest/APIReference/. For example, calls to the <a href="SubmitJob">SubmitJob</a>, <a href="ListJobs">ListJobs</a> and <a href="DescribeJobs">DescribeJobs</a> sections generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or IAM user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another Amazon service.

For more information, see the CloudTrail userIdentity Element.

## Reference: Understanding Amazon Batch log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single

request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the CreateComputeEnvironment action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:admin",
    "arn": "arn:aws-cn:sts::012345678910:assumed-role/Admin/admin",
    "accountId": "012345678910",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2017-12-20T00:48:46Z"
      },
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws-cn:iam::012345678910:role/Admin",
        "accountId": "012345678910",
        "userName": "Admin"
      }
   }
 },
 "eventTime": "2017-12-20T00:48:46Z",
  "eventSource": "batch.amazonaws.com",
  "eventName": "CreateComputeEnvironment",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.1",
  "userAgent": "aws-cli/1.11.167 Python/2.7.10 Darwin/16.7.0 botocore/1.7.25",
  "requestParameters": {
    "computeResources": {
      "subnets": [
        "subnet-5eda8e04"
      ],
      "tags": {
        "testBatchTags": "CLI testing CE"
```

```
"desiredvCpus": 0,
      "minvCpus": 0,
      "instanceTypes": [
        "optimal"
      ],
      "securityGroupIds": [
        "sg-aba9e8db"
      ],
      "instanceRole": "ecsInstanceRole",
      "maxvCpus": 128,
      "type": "EC2"
    },
    "state": "ENABLED",
    "type": "MANAGED",
    "computeEnvironmentName": "Test"
  },
  "responseElements": {
    "computeEnvironmentName": "Test",
    "computeEnvironmentArn": "arn:aws-cn:batch:us-east-1:012345678910:compute-
environment/Test"
  },
  "requestID": "890b8639-e51f-11e7-b038-EXAMPLE",
  "eventID": "874f89fa-70fc-4798-bc00-EXAMPLE",
  "readOnly": false,
  "eventType": "AwsApiCall",
  "recipientAccountId": "012345678910"
}
```

## **Troubleshoot Amazon Batch IAM**

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Batch and IAM.

## **Topics**

- I am not authorized to perform an action in Amazon Batch
- <u>I am not authorized to perform iam:PassRole</u>
- I want to allow people outside of my Amazon account to access my Amazon Batch resources

## I am not authorized to perform an action in Amazon Batch

If the Amazon Web Services Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your user name and password.

The following example error occurs when the mateojackson user tries to use the console to view details about a fictional *my-example-widget* resource but does not have the fictional batch: *GetWidget* permissions.

```
User: arn:aws-cn:iam::123456789012:user/mateojackson is not authorized to perform: batch:GetWidget on resource: my-example-widget
```

In this case, Mateo asks his administrator to update his policies to allow him to access the *my-example-widget* resource using the batch: *GetWidget* action. For more information about granting permissions to pass a role, see <u>Granting a user permissions to pass a role to an Amazon service</u>.

## I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam: PassRole action, your policies must be updated to allow you to pass a role to Amazon Batch.

Some Amazon Web Services services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in Amazon Batch. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws-cn:iam::123456789012:user/marymajor is not authorized to perform: iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the iam: PassRole action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

# I want to allow people outside of my Amazon account to access my Amazon Batch resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Batch supports these features, see How Amazon Batch works with IAM.
- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see <a href="Providing access to an IAM user in another Amazon Web Services account that you own in the IAM User Guide">IAM User Guide</a>.
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see <u>Providing access to Amazon Web Services accounts owned by third parties</u> in the *IAM User Guide*.
- To learn how to provide access through identity federation, see <u>Providing access to externally</u> authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the IAM User Guide.

## **Amazon Step Functions**

You can use the Amazon Batch console to view details about your Step Functions state machines and the functions that they use.

#### **Sections**

- Tutorial: View state machine details
- Tutorial: Edit a state machine
- Tutorial: Run a state machine

## **Tutorial: View state machine details**

The Amazon Batch console displays a list of your state machines in the current Amazon Web Services Region that contain at least one workflow step that submits a Amazon Batch job.

Choose a state machine to view a graphical representation of the workflow. Steps highlighted in blue represent Amazon Batch jobs. Use the graph controls to zoom in, zoom out, and center the graph.



#### Note

When a Amazon Batch job is dynamically referenced with JsonPath in the state machine definition, the function details cannot be shown in the Amazon Batch console. Instead, the job name is listed as a **Dynamic reference**, and the corresponding steps in the graph are grayed out.

#### To view state machine details

- 1. Open the Amazon Batch console Workflow orchestration powered by Step Functions page.
- Choose a state machine.

<result>

The Amazon Batch console opens the **Details** page. </result>

For more information, see Step Functions in the Amazon Step Functions Developer Guide.

## Tutorial: Edit a state machine

When you want to edit a state machine, Amazon Batch opens the **Edit definition** page of the Step Functions console.

#### To edit a state machine

- 1. Open the Amazon Batch console Workflow orchestration powered by Step Functions page.
- 2. Choose a state machine.
- 3. Choose **Edit**.

The Step Functions console opens the **Edit definition** page.

4. Edit the state machine and choose **Save**.

For more information about editing state machines, see <u>Step Functions state machine language</u> in the *Amazon Step Functions Developer Guide*.

## Tutorial: Run a state machine

When you want to run a state machine, Amazon Batch opens the **New execution** page of the Step Functions console.

#### To run a state machine

- 1. Open the Amazon Batch console Workflow orchestration powered by Step Functions page.
- 2. Choose a state machine.
- 3. Choose **Execute**.

The Step Functions console opens the **New execution** page.

4. (Optional) Edit the state machine and choose **Start execution**.

For more information about running state machines, see <u>Step Functions state machine execution</u> concepts in the *Amazon Step Functions Developer Guide*.

Tutorial: Edit a state machine 412

## Amazon Batch event stream for Amazon EventBridge

You can use the Amazon Batch event stream for Amazon EventBridge to receive near real-time notifications regarding the current state of jobs in your job queues.

You can use EventBridge to gain further insights about your Amazon Batch service. More specifically, you can use it to check the progress of jobs, build Amazon Batch custom workflows, generate usage reports or metrics, or build your own dashboards. With Amazon Batch and EventBridge, you don't need scheduling and monitoring code that continuously polls Amazon Batch for job status changes. Instead, you can handle Amazon Batch job state changes asynchronously using a variety of Amazon EventBridge targets. These include Amazon Lambda, Amazon Simple Queue Service, Amazon Simple Notification Service, or Amazon Kinesis Data Streams.

Events from the Amazon Batch event stream are ensured to be delivered at least one time. In the event that duplicate events are sent, the event provides enough information to identify duplicates. That way, you can compare the time stamp of the event and the job status.

Amazon Batch jobs are available as EventBridge targets. Using simple rules, you can match events and submit Amazon Batch jobs in response to them. For more information, see <a href="What is EventBridge">What is EventBridge</a>? in the Amazon EventBridge User Guide. You can also use EventBridge to schedule automated actions that self-trigger at certain times using **cron** or rate expressions. For more information, see <a href="Creating an Amazon EventBridge rule that runs on a schedule">Creating an Amazon EventBridge rule that runs on a schedule</a> in the Amazon EventBridge User Guide. For an example walkthrough, see <a href="Amazon Batch">Amazon Batch</a> jobs as EventBridge targets. For information about using the EventBridge Scheduler, see <a href="Setting up Amazon">Setting up Amazon</a> EventBridge Scheduler in the Amazon EventBridge User Guide.

#### **Topics**

- Amazon Batch events
- Tutorial: Use Amazon user notifications with Amazon Batch
- Amazon Batch jobs as EventBridge targets
- Tutorial: Listen for Amazon Batch job events using EventBridge
- Tutorial: Sending Amazon Simple Notification Service alerts for failed job events

## **Amazon Batch events**

Amazon Batch sends job status change events to EventBridge. Amazon Batch tracks the state of your jobs. If a previously submitted job's status changes, an event is invoked. For example, if a job in the RUNNING status moves to the FAILED status. These events are classified as job state change events.



#### Note

Amazon Batch might add other event types, sources, and details in the future. If you're programmatically deserializing event JSON data, make sure that your application is prepared to handle unknown properties. This is to avoid issues if and when these additional properties are added.

### **Topics**

- Resource: Job state change events
- Resource: Job queue blocked events

## **Resource: Job state change events**

Anytime that an existing (previously submitted) job changes states, an event is created. For more information about Amazon Batch job states, see Job states.



#### Note

Events aren't created for the initial job submission.

## **Example Job State Change Event**

Job state change events are delivered in the following format. The detail section resembles the JobDetail object that's returned from a Describe Jobs API operation in the Amazon Batch API Reference. For more information about EventBridge parameters, see Events and Event Patterns in the Amazon EventBridge User Guide.

```
{
    "version": "0",
```

Amazon Batch events 414

```
"id": "c8f9c4b5-76e5-d76a-f980-7011e206042b",
    "detail-type": "Batch Job State Change",
    "source": "aws.batch",
    "account": "123456789012",
    "time": "2022-01-11T23:36:40Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws-cn:batch:us-east-1:123456789012:job/4c7599ae-0a82-49aa-
ba5a-4727fcce14a8"
    ],
    "detail": {
        "jobArn": "arn:aws-cn:batch:us-east-1:123456789012:job/4c7599ae-0a82-49aa-
ba5a-4727fcce14a8",
        "jobName": "event-test",
        "jobId": "4c7599ae-0a82-49aa-ba5a-4727fcce14a8",
        "jobQueue": "arn:aws-cn:batch:us-east-1:123456789012:job-queue/
PexjEHappyPathCanary2JobQueue",
        "status": "RUNNABLE",
        "attempts": [],
        "createdAt": 1641944200058,
        "retryStrategy": {
            "attempts": 2,
            "evaluateOnExit": []
        },
        "dependsOn": [],
        "jobDefinition": "arn:aws-cn:batch:us-east-1:123456789012:job-definition/first-
run-job-definition:1",
        "parameters": {},
        "container": {
            "image": "137112412989.dkr.ecr.us-east-1.amazonaws.com/amazonlinux:latest",
            "command": [
                "sleep",
                "600"
            ],
            "volumes": [],
            "environment": [],
            "mountPoints": [],
            "ulimits": [],
            "networkInterfaces": [],
            "resourceRequirements": [
                {
                    "value": "2",
                    "type": "VCPU"
                }, {
```

## Resource: Job queue blocked events

Anytime that Amazon Batch detects a job in the RUNNABLE state and thus blocking a queue, an event is created in Amazon CloudWatch Events. For more information about supported blocked queue causes, see <u>Jobs stuck in a RUNNABLE status</u>. The same reason is also available in the statusReason field in the <u>DescribeJobs</u> API action.

## **Example Job State Change Event**

Job state change events are delivered in the following format. The detail section resembles the <u>JobDetail</u> object that's returned from a <u>DescribeJobs</u> API operation in the *Amazon Batch API Reference*. For more information about EventBridge parameters, see <u>Events and Event Patterns</u> in the *Amazon EventBridge User Guide*.

```
"detail": {
        "jobArn": "arn:aws:batch:us-east-1:123456789012:job/4c7599ae-0a82-49aa-
ba5a-4727fcce14a8",
        "jobName": "event-test",
        "jobId": "4c7599ae-0a82-49aa-ba5a-4727fcce14a8",
        "jobQueue": "arn:aws:batch:us-east-1:123456789012:job-queue/
PexjEHappyPathCanary2JobQueue",
        "status": "RUNNABLE",
        "statusReason": "blocked-reason",
        "attempts": [],
        "createdAt": 1641944200058,
        "retryStrategy": {
            "attempts": 2,
            "evaluateOnExit": []
        },
        "depends0n": [],
        "jobDefinition": "arn:aws:batch:us-east-1:123456789012:job-definition/first-
run-job-definition:1",
        "parameters": {},
        "container": {
            "image": "137112412989.dkr.ecr.us-east-1.amazonaws.com/amazonlinux:latest",
            "command": [
                "sleep",
                "600"
            ],
            "volumes": [],
            "environment": [],
            "mountPoints": [],
            "ulimits": [],
            "networkInterfaces": [],
            "resourceRequirements": [
                {
                    "value": "2",
                    "type": "VCPU"
                }, {
                    "value": "256",
                    "type": "MEMORY"
                }
            ],
            "secrets": []
        },
        "tags": {
            "resourceArn": "arn:aws:batch:us-
east-1:123456789012:job/4c7599ae-0a82-49aa-ba5a-4727fcce14a8"
```

```
},
    "propagateTags": false,
    "platformCapabilities": []
}
```

## **Tutorial: Use Amazon user notifications with Amazon Batch**

You can use <u>Amazon User Notifications</u> to set up delivery channels to get notified about Amazon Batch events. You receive a notification when an event matches a rule that you specify. You can receive notifications for events through multiple channels, including email, <u>Amazon Q Developer in chat applications</u> chat notifications, or <u>Amazon Console Mobile Application</u> push notifications. You can also see notifications in the <u>Console Notifications Center</u>. User Notifications supports aggregation, which can reduce the number of notifications you receive during specific events.

To configure User Notifications in Amazon Batch:

- 1. Open the Amazon Batch console.
- 2. Choose Dashboard.
- 3. Choose Configure Notifications.
- 4. In Amazon User Notifications, choose Create notification configuration.

For more information about how to configure and view user notifications, see <u>Getting started with</u> Amazon User Notifications.

## Amazon Batch jobs as EventBridge targets

Amazon EventBridge delivers a near real-time stream of system events that describe changes in Amazon Web Services resources. Typically, Amazon Batch on Amazon Elastic Container Service, Amazon Elastic Kubernetes Service, and Amazon Fargate jobs are available as EventBridge targets. Using simple rules, you can match events and submit Amazon Batch jobs in response to them. For more information, see <a href="What is EventBridge">What is EventBridge</a>? in the Amazon EventBridge User Guide.

You can also use EventBridge to schedule automated actions that are invoked at certain times using **cron** or rate expressions. For more information, see <u>Creating an Amazon EventBridge rule</u> that runs on a schedule in the *Amazon EventBridge User Guide*.

For information about how to create a rule that runs when an event matches an event pattern, see Creating Amazon EventBridge rules that react to events in the Amazon EventBridge User Guide.

Common use cases for Amazon Batch jobs as an EventBridge target include the following use cases:

- A scheduled job occurs at regular time intervals. For example, a cron job occurs only during lowusage hours when Amazon EC2 Spot Instances are less expensive.
- An Amazon Batch job runs in response to an API operation that's logged in CloudTrail. For example, a job is submitted whenever an object is uploaded to a specified Amazon S3 bucket. Each time this happens, the EventBridge input transformer passes the bucket and key name of the object to Amazon Batch parameters.



#### Note

In this scenario, all of related Amazon resources must be in the same Region. This includes resources such as the Amazon S3 bucket, EventBridge rule, and CloudTrail logs.

Before you can submit Amazon Batch jobs with EventBridge rules and targets, the EventBridge service requires several permissions to run Amazon Batch jobs. When you create a rule in the EventBridge console that specifies an Amazon Batch job as a target, you can also create this role. For more information about the required service principal and IAM permissions for this role, see EventBridge IAM role.

#### **Topics**

- Tutorial: Create a scheduled Amazon Batch job
- Tutorial: Create a rule with an event pattern
- Tutorial: Pass event information to an Amazon Batch target on a schedule using the EventBridge input transformer

## **Tutorial: Create a scheduled Amazon Batch job**

The following procedure covers how to create a scheduled Amazon Batch job and the required EventBridge IAM role.

#### To create a scheduled Amazon Batch job with EventBridge



### Note

This procedure works for all Amazon Batch on Amazon ECS, Amazon EKS, and Amazon Fargate jobs.

- Open the Amazon EventBridge console at https://console.amazonaws.cn/events/. 1.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Rules**.
- Choose Create rule. 4.
- For **Name**, specify a unique name for your compute environment. The name can contain up 5. to 64 characters. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).



#### Note

A rule can't have the same name as another rule in the same Region and on the same event bus.

- (Optional) For **Description**, enter a description for the rule. 6.
- 7. For **Event bus**, choose the event bus that you want to associate with this rule. If you want this rule to match events that come from your account, select default. When an Amazon Web Services service in your account emits an event, it always goes to your account's default event bus.
- 8. (Optional) Turn off the rule on the selected bus if you don't want to run the rule immediately.
- 9. For **Rule type**, choose **Schedule**.
- 10. Choose Continue to create rule or Next.
- 11. For **Schedule pattern**, do one of the following:
  - Choose A fine-grained schedule that runs at a specific time, such as 8:00 a.m. PST on the first Monday of every month and then enter a cron expression. For more information, see Cron Expressions in the Amazon EventBridge User Guide.
  - Choose A schedule that runs at a regular rate, such as every 10 minutes. and then enter a rate expression.

- 12. Choose Next.
- 13. For Target types, choose Amazon Web Services service.
- 14. For **Select a target**, choose **Batch job queue**. Then, configure the following:
  - **Job queue:** Enter the Amazon Resource Name (ARN) of the job queue to schedule your job in.
  - **Job definition:** Enter the name and revision or full ARN of the job definition to use for your job.
  - Job name: Enter a name for your job.
  - Array size: (Optional) Enter an array size for your job to run more than one copy. For more information, see Array jobs.
  - **Job attempts:** (Optional) Enter the number of times to retry your job if it fails. For more information, see Automated job retries.
- 15. For **Batch job queue** target types, EventBridge needs permission to send events to the target. EventBridge can create the IAM role needed for your rule to run. Do one of the following:
  - To create an IAM role automatically, choose **Create a new role for this specific resource**.
  - To use an IAM role that you've already created, choose **Use existing role**.
- 16. (Optional) Expand Additional settings.
  - For Configure target input, choose how the text from an event is processed before it's passed to the target.
  - b. For **Maximum age of event**, specify the time interval for how long unprocessed events are kept.
  - c. For **Retry attempts**, enter the number of times that an event is retried.
  - d. For **Dead-letter queue**, choose an option for how unprocessed events are handled. If necessary, specify the Amazon SQS queue to use as the dead-letter queue.
- 17. (Optional) Choose **Add another target** to add another target for this rule.
- 18. Choose **Next**.
- 19. (Optional) For **Tags**, choose **Add new tag** to add a resource label for the rule. For more information, see Amazon EventBridge tags.
- 20. Choose **Next**.
- 21. For **Review and create**, review the configuration steps. If you need to make changes, choose **Edit**. When you're finished, choose **Create rule**.

For more information about creating rules, see Creating an Amazon EventBridge rule that runs on a schedule in the Amazon EventBridge User Guide.

## Tutorial: Create a rule with an event pattern

The following procedure covers how to create a rule with an event pattern.

To create a rule that sends the event to a target when the event matches a defined pattern



This procedure works for all Amazon Batch on Amazon ECS, Amazon EKS, and Amazon Fargate jobs.

- Open the Amazon EventBridge console at https://console.amazonaws.cn/events/. 1.
- From the navigation bar, select the Amazon Web Services Region to use. 2.
- 3. In the navigation pane, choose Rules.
- Choose **Create rule**.
- 5. For **Name**, specify a unique name for your compute environment. The name can contain up to 64 characters. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores (\_).



#### Note

A rule can't have the same name as another rule in the same Region and on the same event bus.

- (Optional) For **Description**, enter a description for the rule. 6.
- 7. For **Event bus**, choose the event bus that you want to associate with this rule. If you want this rule to match events that come from your account, select **default**. When an Amazon Web Services service in your account emits an event, it always goes to your account's default event bus.
- (Optional) Turn off the rule on the selected bus if you don't want to run the rule immediately. 8.
- For **Rule type**, choose **Rule with an event pattern**.
- 10. Choose Next.

- 11. For Event Source, choose Amazon event or EventBridge partner events.
- 12. (Optional) For **Sample event**:
  - a. For **Sample event type**, choose **Amazon events**.
  - b. For Sample events, choose Batch Job State Change.
- 13. For **Creation method**, choose **Use pattern form**.
- 14. For **Event pattern**:
  - a. For Event source, choose Amazon Web Services services.
  - b. For Amazon Web Services service, choose Batch.
  - c. For **Event type**, choose **Batch Job State Change**.
- 15. Choose Next.
- 16. For Target types, choose Amazon Web Services service.
- 17. For Select a target, choose a target type. For example, choose Batch job queue. Then specify the following:
  - **Job queue:** Enter the Amazon Resource Name (ARN) of the job queue to schedule your job in.
  - Job definition: Enter the name and revision or full ARN of the job definition to use for your
    job.
  - Job name: Enter a name for your job.
  - Array size: (Optional) Enter an array size for your job to run more than one copy. For more information, see Array jobs.
  - **Job attempts:** (Optional) Enter the number of times to retry your job if it fails. For more information, see Automated job retries.
- 18. For **Batch job queue** target types, EventBridge needs permission to send events to the target. EventBridge can create the IAM role needed for your rule to run. Do one of the following:
  - To create an IAM role automatically, choose **Create a new role for this specific resource**.
  - To use an IAM role that you created before, choose Use existing role.
- 19. (Optional) Expand Additional settings.
  - a. For **Configure target input**, choose how text from an event is processed.
  - b. For **Maximum age of event**, specify the time interval for how long unprocessed events are kept.

- For **Retry attempts**, enter the number of times that an event is retried. c.
- For **Dead-letter queue**, choose an option for how unprocessed events are handled. If necessary, specify the Amazon SQS queue to use as the dead-letter queue.
- 20. (Optional) Choose **Add another target** to add an additional target.
- 21. Choose Next.
- 22. (Optional) For Tags, choose Add new tag to add a resource label. For more information, see Amazon EventBridge tags in the Amazon EventBridge User Guide.
- 23. Choose Next.
- 24. For **Review and create**, review the configuration steps. If you need to make changes, choose **Edit**. After you're finished, choose **Create rule**.

For more information about creating rules, see Creating Amazon EventBridge rules that react to events in the Amazon EventBridge User Guide.

# Tutorial: Pass event information to an Amazon Batch target on a schedule using the EventBridge input transformer

You can use the EventBridge input transformer to pass event information to Amazon Batch in a job submission. This can be especially valuable if you invoke jobs as a result of other Amazon event information. One example is an object upload to an Amazon S3 bucket. You can also use a job definition with parameter substitution values in the container's command. The EventBridge input transformer can provide the parameter values based on the event data.

Then, afterwards, you create an Amazon Batch event target that parses information from the event that starts it and transforms it into a parameters object. When the job runs, the parameters from the trigger event are passed to the command of the job container.



#### Note

In this scenario, all of the Amazon resources (such as Amazon S3 buckets, EventBridge rules, and CloudTrail logs) must be in the same Region.

#### To create an Amazon Batch target that uses the input transformer

Open the Amazon EventBridge console at https://console.amazonaws.cn/events/.

- From the navigation bar, select the Amazon Web Services Region to use. 2.
- 3. In the navigation pane, choose **Rules**.
- Choose Create rule. 4.
- 5. For **Name**, specify a unique name for your compute environment. The name can contain up to 64 characters. It can contain uppercase and lowercase letters, numbers, hyphens (-), and underscores ().



#### Note

A rule can't have the same name as another rule in the same Amazon Web Services Region and on the same event bus.

- (Optional) For **Description**, enter a description for the rule.
- For **Event bus**, choose the event bus that you want to associate with this rule. If you want this rule to match events that come from your account, select default. When an Amazon Web Services service in your account emits an event, it always goes to your account's default event bus.
- 8. (Optional) Turn off the rule on the selected bus if you don't want to run the rule immediately.
- For **Rule type**, choose **Schedule**. 9.
- Choose Continue to create rule or Next.
- 11. For **Schedule pattern**, do one of the following:
  - Choose A fine-grained schedule that runs at a specific time, such as 8:00 a.m. PST on the first Monday of every month and then enter a cron expression. For more information, see Cron Expressions in the Amazon EventBridge User Guide.
  - Choose A schedule that runs at a regular rate, such as every 10 minutes. and then enter a rate expression.
- 12. Choose Next.
- 13. For Target types, choose Amazon Web Services service.
- 14. For **Select a target**, choose **Batch job queue**. Then, configure the following:
  - Job queue: Enter the Amazon Resource Name (ARN) of the job queue to schedule your job in.
  - Job definition: Enter the name and revision or full ARN of the job definition to use for your job.

- Job name: Enter a name for your job.
- Array size: (Optional) Enter an array size for your job to run more than one copy. For more information, see Array jobs.
- **Job attempts:** (Optional) Enter the number of times to retry your job if it fails. For more information, see <u>Automated job retries</u>.
- 15. For **Batch job queue** target types, EventBridge needs permission to send events to the target. EventBridge can create the IAM role needed for your rule to run. Do one of the following:
  - To create an IAM role automatically, choose Create a new role for this specific resource.
  - To use an IAM role that you've already created, choose **Use existing role**.
- 16. (Optional) Expand Additional settings.
- 17. In the Additional settings section, for Configure target input, choose Input Transformer.
- 18. Choose **Configure input transformer**.
- 19. (Optional) For Sample event:
  - a. For **Sample event type**, choose **Amazon events**.
  - b. For Sample events, choose Batch Job State Change.
- 20. In the **Target input transformer** section, for **Input path**, specify the values to parse from the triggering event. For example, to parse **Batch Job State Change** event, use the following JSON format.

```
{
    "instance": "$.detail.jobId",
    "state": "$.detail.status"
}
```

21. For **Template**, enter the following.

```
{
    "instance": <jobId> ,
    "status": <status>
}
```

- 22. Choose Confirm.
- 23. For **Maximum age of event**, specify the time interval for how long unprocessed events are kept.
- 24. For **Retry attempts**, enter the number of times that an event is retried.

25. For **Dead-letter queue**, choose an option for how unprocessed events are handled. If necessary, specify the Amazon SQS queue to use as the dead-letter queue.

- 26. (Optional) Choose **Add another target** to add an additional target.
- 27. Choose Next.
- 28. (Optional) For **Tags**, choose **Add new tag** to add a resource label. For more information, see <u>Amazon EventBridge tags</u> in the *Amazon EventBridge User Guide*.
- 29. Choose Next.
- 30. For **Review and create**, review the configuration steps. If you need to make changes, choose **Edit**. After you're finished, choose **Create rule**.

# Tutorial: Listen for Amazon Batch job events using EventBridge

In this tutorial, you set up a simple Amazon Lambda function that listens for Amazon Batch job events and writes them out to a CloudWatch Logs log stream.

## **Prerequisites**

This tutorial assumes that you have a working compute environment and job queue that are ready to accept jobs. If you don't have a running compute environment and job queue to capture events from, follow the steps in <u>Getting started with Amazon Batch tutorials</u> to create one. At the end of this tutorial, you can optionally submit a job to this job queue to test that you have configured your Lambda function correctly.

#### **Topics**

- Tutorial: Create the Lambda function
- Tutorial: Register the event rule
- Tutorial: Test your configuration

### **Tutorial: Create the Lambda function**

In this procedure, you create a simple Lambda function to serve as a target for Amazon Batch event stream messages.

#### To create a target Lambda function

1. Open the Amazon Lambda console at https://console.amazonaws.cn/lambda/.

- Choose Create function, Author from scratch. 2.
- 3. For **Function name**, enter **batch-event-stream-handler**.
- For **Runtime**, choose **Python 3.8**. 4.
- Choose **Create function**. 5.
- In the **Code source** section, edit the sample code to match the following example:

```
import json
def lambda_handler(event, _context):
    # _context is not used
    del _context
    if event["source"] != "aws.batch":
        raise ValueError("Function only supports input from events with a source
type of: aws.batch")
    print(json.dumps(event))
```

This is a simple Python 3.8 function that prints the events sent by Amazon Batch. If everything is configured correctly, at the end of this tutorial, the event details appear in the CloudWatch Logs log stream that's associated with this Lambda function.

7. Choose **Deploy**.

## **Tutorial: Register the event rule**

In this section, you create an EventBridge event rule that captures job events that are coming from your Amazon Batch resources. This rule captures all events coming from Amazon Batch within the account where it's defined. The job messages themselves contain information about the event source, including the job queue where it was submitted. You can use this information to filter and sort events programmatically.



#### Note

If you use the Amazon Web Services Management Console to create an event rule, the console automatically adds the IAM permissions for EventBridge to call your Lambda function. However, if you're creating an event rule using the Amazon CLI, you must grant

permissions explicitly. For more information, see <u>Events and Event Patterns</u> in the *Amazon EventBridge User Guide*.

#### To create your EventBridge rule

- 1. Open the Amazon EventBridge console at <a href="https://console.amazonaws.cn/events/">https://console.amazonaws.cn/events/</a>.
- 2. In the navigation pane, choose **Rules**.
- 3. Choose Create rule.
- 4. Enter a name and description for the rule.

A rule can't have the same name as another rule in the same Region and on the same event bus.

- 5. For **Event bus**, choose the event bus that you want to associate with this rule. If you want this rule to match events that come from your account, select **Amazon default event bus**. When an Amazon service in your account emits an event, it always goes to your account's default event bus.
- 6. For Rule type, choose Rule with an event pattern.
- 7. Choose **Next**.
- 8. For **Event source**, choose **Other**.
- 9. For Event pattern, select Custom patterns (JSON editor).
- 10. Paste the following event pattern into the text area.

```
{
   "source": [
     "aws.batch"
]
}
```

This rule applies across all of your Amazon Batch groups and to every Amazon Batch event. Alternatively, you can create a more specific rule to filter out some results.

- 11. Choose Next.
- 12. For **Target types**, choose **Amazon service**.
- 13. For **Select a target**, choose **Lambda function**, and select your Lambda function.
- 14. (Optional) For **Additional settings**, do the following:

a. For **Maximum age of event**, enter a value between one minute (00:01) and 24 hours (24:00).

- b. For **Retry attempts**, enter a number between 0 and 185.
- c. For **Dead-letter queue**, choose whether to use a standard Amazon SQS queue as a dead-letter queue. EventBridge sends events that match this rule to the dead-letter queue if they are not successfully delivered to the target. Do one of the following:
  - Choose **None** to not use a dead-letter queue.
  - Choose Select an Amazon SQS queue in the current Amazon account to use as the dead-letter queue and then select the queue to use from the dropdown.
  - Choose Select an Amazon SQS queue in an other Amazon account as a dead-letter
    queue and then enter the ARN of the queue to use. You must attach a resource-based
    policy to the queue that grants EventBridge permission to send messages to it. For
    more information, see <u>Granting permissions to the dead-letter queue</u> in the *Amazon*EventBridge User Guide.
- 15. Choose Next.
- 16. (Optional) Enter one or more tags for the rule. For more information, see <u>Amazon EventBridge</u> <u>tags</u> in the *Amazon EventBridge User Guide*.
- 17. Choose Next.
- 18. Review the details of the rule and choose **Create rule**.

# **Tutorial: Test your configuration**

You can now test your EventBridge configuration by submitting a job to your job queue. If everything is configured properly, your Lambda function is triggered and it writes the event data to a CloudWatch Logs log stream for the function.

#### To test your configuration

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. Submit a new Amazon Batch job. For more information, see <u>Tutorial: submit a job</u>.
- 3. Open the CloudWatch console at <a href="https://console.amazonaws.cn/cloudwatch/">https://console.amazonaws.cn/cloudwatch/</a>.
- 4. On the navigation pane, choose **Logs** and select the log group for your Lambda function (for example, /aws/lambda/my-function).
- 5. Select a log stream to view the event data.

# Tutorial: Sending Amazon Simple Notification Service alerts for failed job events

In this tutorial, you configure an Amazon EventBridge event rule that only captures job events where the job has moved to a FAILED status. At the end of this tutorial, you can optionally also submit a job to this job queue. This is to test that you have configured your Amazon SNS alerts correctly.

# **Prerequisites**

This tutorial assumes that you have a working compute environment and job queue that are ready to accept jobs. If you don't have a running compute environment and job queue to capture events from, follow the steps in Getting started with Amazon Batch tutorials to create one.

#### **Topics**

- Tutorial: Create and subscribe to an Amazon SNS topic
- Tutorial: Register an event rule
- Tutorial: Test your rule
- Alternate rule: Batch job queue blocked

# **Tutorial: Create and subscribe to an Amazon SNS topic**

For this tutorial, you configure an Amazon SNS topic to serve as an event target for your new event rule.

#### To create an Amazon SNS topic

- 1. Open the Amazon SNS console at https://console.amazonaws.cn/sns/v3/home.
- 2. Choose **Topics**, **Create topic**.
- 3. For **Type**, choose **Standard**.
- 4. For Name, enter JobFailedAlert and choose Create topic.
- 5. On the **JobFailedAlert** screen, choose **Create subscription**.
- 6. For **Protocol**, choose **Email**.
- 7. For **Endpoint**, enter an email address that you currently have access to and choose **Create** subscription.

8. Check your email account, and wait to receive a subscription confirmation email message. When you receive it, choose **Confirm subscription**.

# Tutorial: Register an event rule

Next, register an event rule that captures only job-failed events.

#### To register your EventBridge rule

- 1. Open the Amazon EventBridge console at <a href="https://console.amazonaws.cn/events/">https://console.amazonaws.cn/events/</a>.
- 2. In the navigation pane, choose **Rules**.
- 3. Choose Create rule.
- 4. Enter a name and description for the rule.

A rule can't have the same name as another rule in the same Region and on the same event bus.

- 5. For **Event bus**, choose the event bus that you want to associate with this rule. If you want this rule to match events that come from your account, select **Amazon default event bus**. When an Amazon service in your account emits an event, it always goes to your account's default event bus.
- 6. For Rule type, choose Rule with an event pattern.
- 7. Choose Next.
- 8. For **Event source**, choose **Other**.
- 9. For **Event pattern**, select **Custom patterns (JSON editor)**.
- 10. Paste the following event pattern into the text area.

```
{
  "detail-type": [
    "Batch Job State Change"
],
  "source": [
    "aws.batch"
],
  "detail": {
    "status": [
    "FAILED"
]
```

}

This code defines an EventBridge rule that matches any event where the job status is FAILED. For more information about event patterns, see <a href="Events and Event Patterns">Events and Event Patterns</a> in the Amazon EventBridge User Guide.

- 11. Choose Next.
- 12. For **Target types**, choose **Amazon service**.
- 13. For Select a target, choose SNS topic, and for Topic, choose JobFailedAlert.
- 14. (Optional) For **Additional settings**, do the following:
  - a. For **Maximum age of event**, enter a value between one minute (00:01) and 24 hours (24:00).
  - b. For **Retry attempts**, enter a number between 0 and 185.
  - c. For **Dead-letter queue**, choose whether to use a standard Amazon SQS queue as a dead-letter queue. EventBridge sends events that match this rule to the dead-letter queue if they are not successfully delivered to the target. Do one of the following:
    - Choose None to not use a dead-letter queue.
    - Choose Select an Amazon SQS queue in the current Amazon account to use as the dead-letter queue and then select the queue to use from the dropdown.
    - Choose Select an Amazon SQS queue in an other Amazon account as a dead-letter
      queue and then enter the ARN of the queue to use. You must attach a resource-based
      policy to the queue that grants EventBridge permission to send messages to it. For
      more information, see <u>Granting permissions to the dead-letter queue</u> in the *Amazon*EventBridge User Guide.
- 15. Choose Next.
- 16. (Optional) Enter one or more tags for the rule. For more information, see <u>Amazon EventBridge</u> tags in the *Amazon EventBridge User Guide*.
- 17. Choose Next.
- 18. Review the details of the rule and choose **Create rule**.

# **Tutorial: Test your rule**

To test your rule, submit a job that exits shortly after it starts with a non-zero exit code. If your event rule is configured correctly, you should receive an email message within a few minutes with the event text.

#### To test a rule

- Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. Submit a new Amazon Batch job. For more information, see <u>Tutorial</u>: <u>submit a job</u>. For the job's command, substitute this command to exit the container with an exit code of 1.

```
/bin/sh, -c, 'exit 1'
```

3. Check your email to confirm that you received an email alert for the failed job notification.

## Alternate rule: Batch job queue blocked

To create an event rule that monitors for *batch job queue blocked*, repeat these tutorials with the following alterations:

- 1. In <u>Tutorial: Create and subscribe to an Amazon SNS topic</u>, use <u>BlockedJobQueue</u> as the topic name.
- 2. In Tutorial: Register an event rule, use the following pattern in the JSON editor:

```
{
  "detail-type": [
     "Batch Job Queue Blocked"
],
  "source": [
     "aws.batch"
]
}
```

Tutorial: Test your rule 434

# **Elastic Fabric Adapter**

An Elastic Fabric Adapter (EFA) is a network device to accelerate High Performance Computing (HPC) applications. Amazon Batch supports applications that use EFA if the following conditions are met.

 For a list of instance types that support EFAs, see Supported instance types in the Amazon EC2 User Guide.



#### (i) Tip

To see a list of instance types that support EFAs in an Amazon Web Services Region, run the following command. Then, cross reference the list that's returned with the list of available instance types in the Amazon Batch console.

```
$ aws ec2 describe-instance-types --region us-east-1 --filters Name=network-
info.efa-supported,Values=true --query "InstanceTypes[*].[InstanceType]"
output text | sort
```

- For a list of operating systems that support EFA, see Supported operating systems.
- The AMI has the EFA driver loaded.
- The security group for the EFA must allows all inbound and outbound traffic to and from the security group itself.
- All instances that use an EFA must be in the same cluster placement group.
- The job definition must include a devices member with hostPath set to /dev/infiniband/ uverbs0 to allow the EFA device to be passed through to the container. If containerPath is specified, it must also be set to /dev/infiniband/uverbs0. If permissions is set it must be set to READ | WRITE | MKNOD.

The location of the LinuxParameters members are different for multi-node parallel jobs and single-node container jobs. The following examples show the differences, but are missing required values.

#### Example Example for multi-node parallel job

```
"jobDefinitionName": "EFA-MNP-JobDef",
```

```
"type": "multinode",
  "nodeProperties": {
    "nodeRangeProperties": [
      {
        "container": {
          . . .
          "linuxParameters": {
            "devices": [
              {
                 "hostPath": "/dev/infiniband/uverbs0",
                 "containerPath": "/dev/infiniband/uverbs0",
                 "permissions": [
                     "READ", "WRITE", "MKNOD"
              },
            ],
          },
        },
      },
   ],
  },
}
```

### Example Example for single-node container job

For more information about EFA, see <u>Elastic Fabric Adapter</u> in *Amazon EC2 User Guide*.

## **Monitor Amazon Batch**

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon Batch and your Amazon solution.

We strongly encourage you to collect monitoring data from all parts of your Amazon solution to make it easier to debug a multi-point failure, if one occurs. Start by creating a monitoring plan that answers the following questions. If you're not sure how to answer these, you can still use Amazon CloudWatch Logs to establish your performance baselines.

- What are your monitoring goals?
- Which resources will you monitor?
- How often will you monitor these resources?
- Which monitoring tools will you use?
- Who will perform the monitoring tasks?
- Who should be notified when something goes wrong?

Your next step is to establish a baseline of normal Amazon Batch performance in your environment by measuring performance at various times and under different load conditions. As you monitor Amazon Batch, keep historical monitoring data so that you can compare it with current performance data. This will help you identify normal performance patterns and performance anomalies, and devise methods to address issues.

The topics in this section can help you start logging and monitoring Amazon Batch.

#### **Topics**

- Using CloudWatch Logs with Amazon Batch
- Amazon Batch CloudWatch Container Insights
- Use CloudWatch Logs to monitor Amazon Batch on Amazon EKS jobs

# **Using CloudWatch Logs with Amazon Batch**

You can configure your Amazon Batch jobs on EC2 resources to send detailed log information and metrics to CloudWatch Logs. Doing this, you can view different logs from your jobs in

CloudWatch Logs 438

one convenient location. For more information about CloudWatch Logs, see What is Amazon CloudWatch Logs? in the Amazon CloudWatch User Guide.



#### Note

By default, CloudWatch Logs is turned on for Amazon Fargate containers.

To turn on and customize CloudWatch Logs logging, review the following one-time configuration tasks:

- For Amazon Batch compute environments that are based on EC2 resources, add an IAM policy to the ecsInstanceRole role. For more information, see the section called "Tutorial: Add a CloudWatch Logs IAM policy".
- Create an Amazon EC2 launch template that includes detailed CloudWatch monitoring, then specify the template when you create your Amazon Batch compute environment. You can also install the CloudWatch agent on an existing image and then specify the image in the Amazon Batch first-run wizard.
- (Optional) Configure the awslogs driver. You can add parameters that change the default behavior on both EC2 and Fargate resources. For more information, see the section called "Use the awslogs log driver".

#### **Topics**

- Tutorial: Add a CloudWatch Logs IAM policy
- Install and configure the CloudWatch agent
- Tutorial: View CloudWatch Logs

## **Tutorial: Add a CloudWatch Logs IAM policy**

Before your jobs can send log data and detailed metrics to CloudWatch Logs, you must create an IAM policy that uses the CloudWatch Logs APIs. After you create the IAM policy, attach it to the ecsInstanceRole role.



#### Note

If the ECS-CloudWatchLogs policy isn't attached to the ecsInstanceRole role, basic metrics can still be sent to CloudWatch Logs. However, the basic metrics don't include log data or detailed metrics such as free disk space.

Amazon Batch compute environments use Amazon EC2 resources. When you create a compute environment using the Amazon Batch first-run wizard, Amazon Batch creates the ecsInstanceRole role and configures the environment with it.

If you aren't using the first-run wizard, you can specify the ecsInstanceRole role when you create a compute environment in the Amazon Command Line Interface or Amazon Batch API. For more information, see the Amazon CLI Command Reference or Amazon Batch API Reference.

#### To create the ECS-CloudWatchLogs IAM policy

- 1. Open the IAM console at https://console.amazonaws.cn/iam/.
- 2. In the navigation pane, choose **Policies**.
- 3. Choose **Create policy**.
- 4. Choose **JSON**, then enter the following policy:

**JSON** 

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogGroup",
                "logs:CreateLogStream",
                "logs:PutLogEvents",
                "logs:DescribeLogStreams"
            ],
            "Resource": [
                "arn:aws-cn:logs:*:*:*"
            ]
```

> ] }

- Choose Next: Tags. 5.
- (Optional) For **Add tags**, choose **Add tag** to add a tag to the policy.
- 7. Choose Next: Review.
- On the **Review policy** page, for **Name**, enter **ECS-CloudWatchLogs**, and then enter an optional **Description**.
- Choose **Create policy**. 9.

#### To attach the ECS-CloudWatchLogs policy to ecsInstanceRole

- Open the IAM console at https://console.amazonaws.cn/iam/.
- 2. In the navigation pane, choose **Roles**.
- Choose ecsInstanceRole. If the role doesn't exist, follow the procedures in Amazon ECS instance role to create the role.
- Choose **Add Permissions**, then choose **Attach policies**. 4.
- 5. Choose the **ECS-CloudWatchLogs** policy and then choose **Attach policy**.

## Install and configure the CloudWatch agent

You can create an Amazon EC2 launch template that includes CloudWatch monitoring. For more information, see Launch an instance from a launch template and Advanced details in the Amazon EC2 User Guide.

You can also install the CloudWatch agent on an existing Amazon EC2 AMI and then specify the image in the Amazon Batch first-run wizard. For more information, see Installing the CloudWatch agent and Getting started with Amazon Batch tutorials.



#### Note

Launch templates are not supported on Amazon Fargate resources.

## **Tutorial: View CloudWatch Logs**

You can view and search CloudWatch Logs logs in the Amazon Web Services Management Console.

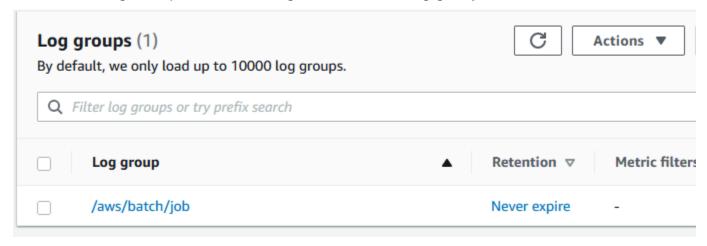


#### Note

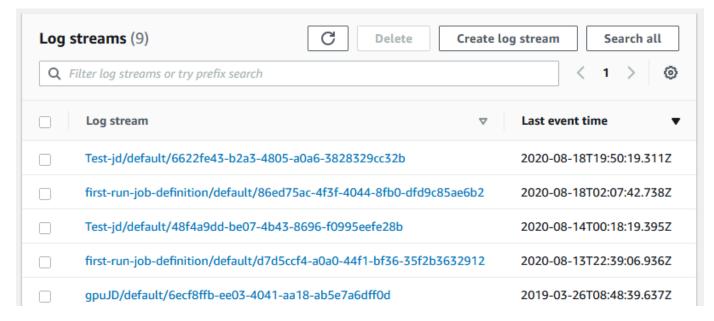
It might take a few minutes for data to display in CloudWatch Logs.

#### To view your CloudWatch Logs data

- Open the CloudWatch console at https://console.amazonaws.cn/cloudwatch/. 1.
- In the left navigation pane, choose **Logs**, then choose **Log groups**. 2.



3. Choose a log group to view.

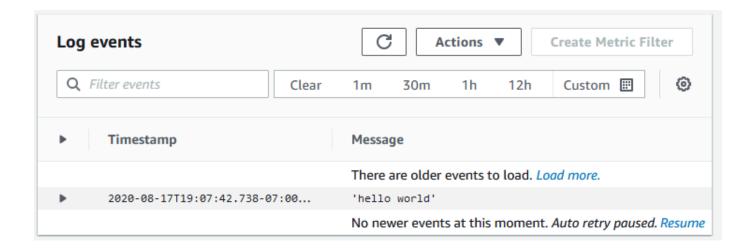


Choose a log stream to view. By default, the streams are identified by the first 200 characters 4. of the job name and the Amazon ECS task ID.



#### (i) Tip

To download log stream data, choose **Actions**.



# **Amazon Batch CloudWatch Container Insights**

CloudWatch Container Insights collects, aggregates, and summarizes metrics and logs from your Amazon Batch compute environments and jobs. The metrics include CPU, memory, disk, and network utilization. You can add these metrics to CloudWatch dashboards.

Operational data is collected as performance log events. These are entries that use a structured JSON schema that enables high-cardinality data to be ingested and stored at scale. From this data, CloudWatch creates higher-level aggregated metrics at the compute environment and job level as CloudWatch metrics. For more information, see Container Insights Structured Logs for Amazon ECS in the Amazon CloudWatch User Guide.



#### Important

CloudWatch Container Insights are charged as custom metrics by CloudWatch. For more information, see Amazon CloudWatch Events pricing

#### **Topics**

Tutorial: Turn on Container Insights

# **Tutorial: Turn on Container Insights**

Complete the following steps to turn on Container Insights for Amazon Batch compute environments.

- 1. Open the Amazon Batch console.
- 2. Choose Compute Environments.
- 3. Choose the compute environment that you want.
- 4. For **Container Insights**, turn on Container Insights for the compute environment.



#### (i) Tip

You can select a default interval to aggregate the metrics or create a custom interval.

By default, the following metrics are displayed. For a full list of Amazon ECS Container Insights metrics, see Amazon ECS Container Insights Metrics in the Amazon CloudWatch User Guide.

- **JobCount** The number of jobs that run in the compute environment.
- ContainerInstanceCount The number of Amazon Elastic Compute Cloud instances that run the Amazon ECS agent and are registered in the compute environment.
- MemoryReserved The memory that's reserved by compute environment jobs. This metric is collected only for the jobs that have a defined memory reservation in their job definition.
- MemoryUtilized The memory that's being used by compute environment jobs. This metric is collected only for jobs that have a defined memory reservation in their job definition.
- CpuReserved The CPU units that are reserved by compute environment jobs. This metric is collected only for jobs that have a defined CPU reservation in their job definition.
- CpuUtilized The CPU units used by jobs in the compute environment. This metric is collected only for jobs that have a defined CPU reservation in their job definition.
- NetworkRxBytes The number of bytes that are received. This metric is available only for containers in jobs that use the awsvpc or bridge network modes.
- NetworkTxBytes The number of bytes that are transmitted. This metric is available only for containers in jobs that use the awsvpc or bridge network modes.
- **StorageReadBytes** The number of bytes that are read from storage.
- **StorageWriteBytes** The number of bytes that are written to storage.

# Use CloudWatch Logs to monitor Amazon Batch on Amazon EKS jobs

You can use Amazon CloudWatch Logs to monitor, store, and view all your log files in one location. Using CloudWatch Logs, you can search, filter, and analyze log data from multiple sources.

You can download an Amazon for Fluent Bit image that includes a plugin to monitor Amazon Batch on Amazon EKS jobs in CloudWatch Logs. Fluent Bit is an open-source log processor and forwarder that's both Docker and Kubernetes compatible. We recommend that you use Fluent Bit as your log router because it's less resource intensive than Fluentd. For more information, see <a href="Install the CloudWatch agent with the Amazon CloudWatch Observability EKS add-on or the Helm chart.">Install the CloudWatch agent with the Amazon CloudWatch Observability EKS add-on or the Helm chart.</a>

## **Prerequisites**

 Attach the CloudWatchAgentServerPolicy policy to the Amazon Identity and Access Management policy of your worker nodes. For more information, see Verify prerequisites.

#### Install the add-on

For instructions on how to install Amazon for Fluent Bit and create the CloudWatch groups, see <a href="Install the CloudWatch agent with the Amazon CloudWatch Observability EKS add-on or the Helm chart">Install the CloudWatch agent with the Amazon CloudWatch Observability EKS add-on or the Helm chart</a>.

You must provide the following additional configuration data when installing the add-on:

• If you install the add-on with the Amazon Web Services Management Console you need to provide the following tolerations in **Configuration values**:

```
{
  "tolerations": [
    {
        "key": "batch.amazonaws.com/batch-node",
        "operator": "Exists"
    }
  ]
}
```

• If you install the add-on with the Amazon CLI then add the following arguments:

--configuration-values '{"tolerations":[{"key":"batch.amazonaws.com/batchnode","operator":"Exists"}]}'



# (i) Tip

Remember that Fluent Bit uses .5 CPU and 100 MB of memory on Amazon Batch nodes. This reduces the total available capacity for Amazon Batch jobs. Consider this when you size your jobs.

Install the add-on 446

# Tag your Amazon Batch resources

To help you manage your Amazon Batch resources, you can assign your own metadata to each resource in the form of *tags*. This topic describes tags and shows you how to create them.

#### **Topics**

- Tag basics
- Tag your resources
- Tag restrictions
- Tutorial: Manage tags using the console
- Manage tags using the CLI or API

# Tag basics

A tag is a label that you assign to an Amazon resource. Each tag consists of a *key* and an optional *value*, both of which you define.

Tags enable you to categorize your Amazon resources by, for example, purpose, owner, or environment. When you have many resources of the same type, you can quickly identify a specific resource based on the tags you've assigned to it. For example, you can define a set of tags for your Amazon Batch services to help you track each service's owner and stack level. We recommend that you devise a consistent set of tag keys for each resource type.

Tags are not automatically assigned to your resources. After you add a tag, you can edit tag keys and values or remove tags from a resource at any time. If you delete a resource, any tags for the resource are also deleted.

Tags don't have any semantic meaning to Amazon Batch and are interpreted strictly as a string of characters. You can set the value of a tag to an empty string, but you can't set the value of a tag to null. If you add a tag that has the same key as an existing tag on that resource, the new value overwrites the old value.

You can work with tags using the Amazon Web Services Management Console, the Amazon CLI, and the Amazon Batch API.

If you're using Amazon Identity and Access Management (IAM), you can control which users in your Amazon account have permission to create, edit, or delete tags.

Tag basics 447

## Tag your resources

You can tag new or existing Amazon Batch compute environments, jobs, job definitions, job queues, and scheduling policies.

If you're using the Amazon Batch console, you can apply tags to new resources when they are created or to existing resources at any time using the **Tags** tab on the relevant resource page.

If you're using the Amazon Batch API, the Amazon CLI, or an Amazon SDK, you can apply tags to new resources using the tags parameter on the relevant API action or to existing resources using the TagResource API action. For more information, see <u>TagResource</u>.

Some resource-creating actions enable you to specify tags for a resource when the resource is created. If tags cannot be applied during resource creation, the resource creation process fails. This ensures that resources you intended to tag on creation are either created with specified tags or not created at all. If you tag resources at the time of creation, you don't need to run custom tagging scripts after resource creation.

The following table describes the Amazon Batch resources that can be tagged, and the resources that can be tagged on creation.

#### Tag support for Amazon Batch resources

Resource	Supports tags	Supports tag propagation	Supports tagging on creation (Amazon Batch API, Amazon CLI, Amazon SDK)
Amazon Batch compute environme nts	Yes	No. Compute environment tags do not propagate to any other resources. Tags for the resources are specified in the tags member of the computeRe sources object passed in the <u>CreateCom</u>	Yes

Tag your resources 448

Resource	Supports tags	Supports tag propagation	Supports tagging on creation (Amazon Batch API, Amazon CLI, Amazon SDK)
		<u>puteEnvironment</u> API operation.	
Amazon Batch jobs	Yes	Yes	Yes
Amazon Batch job definitions	Yes	No	Yes
Amazon Batch job queues	Yes	No	Yes
Amazon Batch scheduling policies	Yes	No	Yes

# Tag restrictions

The following basic restrictions apply to tags:

- Maximum number of tags per resource 50
- For each resource, each tag key must be unique, and each tag key can have only one value.
- Maximum key length 128 Unicode characters in UTF-8
- Maximum value length 256 Unicode characters in UTF-8
- If your tagging schema is used across multiple Amazon services and resources, remember that other services may have restrictions on allowed characters. Generally allowed characters are letters, numbers, spaces representable in UTF-8, and the following characters: + = . \_ : / @.
- Tag keys and values are case sensitive.
- Don't use aws:, AWS:, or any upper or lowercase combination of such as a prefix for either keys or values, as it is reserved for Amazon use. You can't edit or delete tag keys or values with this prefix. Tags with this prefix do not count against your tags-per-resource limit.

Tag restrictions 449

# Tutorial: Manage tags using the console

Using the Amazon Batch console, you can manage the tags associated with new or existing compute environments, jobs, job definitions, and job queues.

## Add tags on an individual resource on creation

You can add tags to Amazon Batch compute environments, jobs, job definitions, job queues, and scheduling policies when you create them.

## Add and delete tags on an individual resource

Amazon Batch allows you to add or delete tags associated with your clusters directly from the resource's page.

#### To add or delete a tag on an individual resource

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, choose the Region to use.
- 3. In the navigation pane, choose a resource type (for example, **Job Queues**).
- 4. Choose a specific resource, then choose **Edit tags**.
- 5. Add or delete your tags as necessary.
  - To add a tag specify the key and value in the empty text boxes at the end of the list.
  - To delete a tag choose the

Delete icon button next to the tag.

6. Repeat this process for each tag you want to add or delete, and then choose **Edit tags** to finish.

# Manage tags using the CLI or API

Use the following Amazon CLI commands or Amazon Batch API operations to add, update, list, and delete the tags for your resources.

#### Tag support for Amazon Batch resources

Task	API action	Amazon CLI	Amazon Tools for Windows PowerShell
Add or overwrite one or more tags.	<u>TagResource</u>	tag-resource	Add-BATResourceTag
Delete one or more tags.	UntagResource	untag-resource	Remove-BATResourceTag
List tags for a resource	<u>ListTagsForResourc</u> <u>e</u>	list-tags-for-reso urce	Get-BATResourceTag

The following examples show how to tag or untag resources using the Amazon CLI.

#### **Example 1: Tag an existing resource**

The following command tags an existing resource.

```
aws batch tag-resource --resource-arn resource_ARN --tags team=devs
```

#### Example 2: Untag an existing resource

The following command deletes a tag from an existing resource.

```
aws batch untag-resource --resource-arn resource_ARN --tag-keys tag_key
```

#### **Example 3: List tags for a resource**

The following command lists the tags associated with an existing resource.

```
aws batch list-tags-for-resource --resource-arn resource_ARN
```

Some resource-creating actions enable you to specify tags when you create the resource. The following actions support tagging on creation.

Task	API action	Amazon CLI	Amazon Tools for Windows PowerShell
Create a compute environment	CreateComputeEnvir onment	create-compute-env ironment	New-BATComputeEnvi ronment
Create a job queue	CreateJobQueue	create-job-queue	New-BATJobQueue
Create a scheduling policy	CreateSchedulingPo licy	create-scheduling- policy	New-BATSchedulingP olicy
Register a job definition	RegisterJobDefinition	register-job-defin ition	Register-BATJobDef inition
Submit a job	SubmitJob	submit-job	Submit-BATJob
Create a consumable resource	<u>CreateConsumableRe</u> <u>source</u>	create-consumable- resource	Create-BATConsumab leResource

# **Best practices for Amazon Batch**

You can use Amazon Batch to run a variety of demanding computational workloads at scale without managing a complex architecture. Amazon Batch jobs can be used in a wide range of use cases in areas such as epidemiology, gaming, and machine learning.

This topic covers the best practices to consider while using Amazon Batch and guidance on how to run and optimize your workloads when using Amazon Batch.

#### **Topics**

- When to use Amazon Batch
- Checklist to run at scale
- Optimize containers and AMIs
- Choose the right compute environment resource
- Amazon EC2 On-Demand or Amazon EC2 Spot
- Use Amazon EC2 Spot best practices for Amazon Batch
- Common errors and troubleshooting

## When to use Amazon Batch

Amazon Batch runs jobs at scale and at low cost, and provides queuing services and cost-optimized scaling. However, not every workload is suitable to be run using Amazon Batch.

- Short jobs If a job runs for only a few seconds, the overhead to schedule the batch job might take longer than the runtime of the job itself. As a workaround, binpack your tasks together before you submit them in Amazon Batch. Then, configure your Amazon Batch jobs to iterate over the tasks. For example, stage the individual task arguments into an Amazon DynamoDB table or as a file in an Amazon S3 bucket. Consider grouping tasks so the jobs run 3-5 minutes each. After you binpack the jobs, loop through your task groups within your Amazon Batch job.
- Jobs that must be run immediately Amazon Batch can process jobs quickly. However, Amazon Batch is a scheduler and optimizes for cost performance, job priority, and throughput. Amazon Batch might require time to process your requests. If you need a response in under a few seconds, then a service-based approach using Amazon ECS or Amazon EKS is more suitable.

When to use Amazon Batch 453

## Checklist to run at scale

Before you run a large workload on 50 thousand or more vCPUs, consider the following checklist.



#### Note

If you plan to run a large workload on a million or more vCPUs or need guidance running at large scale, contact your Amazon team.

- Check your Amazon EC2 quotas Check your Amazon EC2 quotas (also known as limits) in the Service Quotas panel of the Amazon Web Services Management Console. If necessary, request a quota increase for your peak number of Amazon EC2 instances. Remember that Amazon EC2 Spot and Amazon On-Demand instances have separate quotas. For more information, see Getting started with Service Quotas.
- Verify your Amazon Elastic Block Store quota for each Region Each instance uses a GP2 or GP3 volume for the operating system. By default, the quota for each Amazon Web Services Region is 300 TiB. However, each instance uses counts as part of this quota. So, make sure to factor this in when you verify your Amazon Elastic Block Store quota for each Region. If your quota is reached, you can't create more instances. For more information, see Amazon Elastic Block Store endpoints and quotas
- Use Amazon S3 for storage Amazon S3 provides high throughput and helps to eliminate the guesswork on how much storage to provision based on the number of jobs and instances in each Availability Zone. For more information, see Best practices design patterns: optimizing Amazon S3 performance.
- Scale gradually to identify bottlenecks early For a job that runs on a million or more vCPUs, start lower and gradually increase so that you can identify bottlenecks early. For example, start by running on 50 thousand vCPUs. Then, increase the count to 200 thousand vCPUs, and then 500 thousand vCPUs, and so on. In other words, continue to gradually increase the vCPU count until you reach the desired number of vCPUs.
- Monitor to identify potential issues early To avoid potential breaks and issues when running at scale, make sure to monitor both your application and architecture. Breaks might occur even when scaling from 1 thousand to 5 thousand vCPUs. You can use Amazon CloudWatch Logs to review log data or use CloudWatch Embedded Metrics using a client library. For more information, see CloudWatch Logs agent reference and aws-embedded-metrics

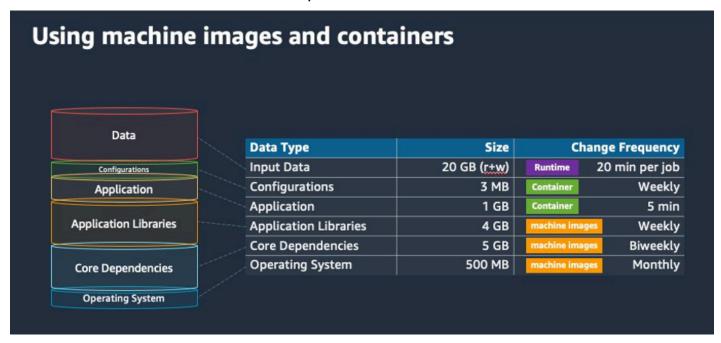
Checklist to run at scale 454

# **Optimize containers and AMIs**

Container size and structure are important for the first set of jobs that you run. This is especially true if the container is larger than 4 GB. Container images are built in layers. The layers are retrieved in parallel by Docker using three concurrent threads. You can increase the number of concurrent threads using the max-concurrent-downloads parameter. For more information, see the Dockerd documentation.

Although you can use larger containers, we recommend that you optimize container structure and size for faster startup times.

- Smaller containers are fetched faster Smaller containers can lead to faster application start times. To decrease container size, offload libraries or files that are updated infrequently to the Amazon Machine Image (AMI). You can also use bind mounts to give access to your containers. For more information, see Bind mounts.
- Create layers that are even in size and break up large layers Each layer is retrieved by one thread. So, a large layer might significantly impact your job startup time. We recommend a maximum layer size of 2 GB as a good tradeoff between larger container size and faster startup times. You can run the docker history your\_image\_id command to check your container image structure and layer size. For more information, see the Docker documentation.
- Use Amazon Elastic Container Registry as your container repository When you run thousands of jobs in parallel, a self-managed repository can fail or throttle throughput. Amazon ECR works at scale and can handle workloads with up to over a million vCPUs.



Optimize containers and AMIs 455

# Choose the right compute environment resource

Amazon Fargate requires less initial setup and configuration than Amazon EC2 and is likely easier to use, particularly if it's your first time. With Fargate, you don't need to manage servers, handle capacity planning, or isolate container workloads for security.

If you have the following requirements, we recommend you use Fargate instances:

- Your jobs must start quickly, specifically less than 30 seconds.
- The requirements of your jobs are 16 vCPUs or less, no GPUs, and 120 GiB of memory or less.

For more information, see When to use Fargate.

If you have the following requirements, we recommend that you use Amazon EC2 instances:

- You require increased control over the instance selection or require using specific instance types.
- Your jobs require resources that Amazon Fargate can't provide, such as GPUs, more memory, a custom AMI, or the Amazon Elastic Fabric Adapter.
- You require a high level of throughput or concurrency.
- You need to customize your AMI, Amazon EC2 Launch Template, or access to special Linux parameters.

With Amazon EC2, you can more finely tune your workload to your specific requirements and run at scale if needed.

# **Amazon EC2 On-Demand or Amazon EC2 Spot**

Most Amazon Batch customers use Amazon EC2 Spot instances because of the savings over On-Demand instances. However, if your workload runs for multiple hours and can't be interrupted, On-Demand instances might be more suitable for you. You can always try Spot instances first and switch to On-Demand if necessary.

If you have the following requirements and expectations, use Amazon EC2 On-Demand instances:

 The runtime of your jobs is more than an hour, and you can't tolerate interruptions to your workload.

 You have a strict SLO (service-level objective) for your overall workload and can't increase computational time.

The instances that you require are more likely to see interruptions.

If you have the following requirements and expectations, use Amazon EC2 Spot instances:

- The runtime for your jobs is typically 30 minutes or less.
- You can tolerate potential interruptions and job rescheduling as a part of your workload. For more information, see Spot Instance advisor.
- Long running jobs can be restarted from a checkpoint if interrupted.

You can mix both purchasing models by submitting on Spot instance first and then use On-Demand instance as a fallback option. For example, submit your jobs on a queue that's connected to compute environments that are running on Amazon EC2 Spot instances. If a job gets interrupted, catch the event from Amazon EventBridge and correlate it to a Spot instance reclamation. Then, resubmit the job to an On-Demand queue using an Amazon Lambda function or Amazon Step Functions. For more information, see Tutorial: Sending Amazon Simple Notification Service alerts for failed job events, Best practices for handling Amazon EC2 Spot Instance interruptions and Manage Amazon Batch with Step Functions.

#### Important

Use different instance types, sizes, and Availability Zones for your On-Demand compute environment to maintain Amazon EC2 Spot instance pool availability and decrease the interruption rate.

## Use Amazon EC2 Spot best practices for Amazon Batch

When you choose Amazon Elastic Compute Cloud (EC2) Spot instances, you likely can optimize your workflow to save costs, sometimes significantly. For more information, see Best practices for Amazon EC2 Spot.

To optimize your workflow to save costs, consider the following Amazon EC2 Spot best practices for Amazon Batch:

 Choose the SPOT\_CAPACITY\_OPTIMIZED allocation strategy – Amazon Batch chooses Amazon EC2 instances from the deepest Amazon EC2 Spot capacity pools. If you're concerned about interruptions, this is a suitable choice. For more information, see <u>Instance type allocation</u> strategies for Amazon Batch.

- Diversify instance types To diversify your instance types, consider compatible sizes and families, then let Amazon Batch choose based on price or availability. For example, consider c5.24xlarge as an alternative to c5.12xlarge or c5a, c5n, c5d, m5, and m5d families. For more information, see Be flexible about instance types and Availability Zones.
- Reduce job runtime or checkpoint We advise against running jobs that take an hour or more when using Amazon EC2 Spot instances to avoid interruptions. If you divide or checkpoint your jobs into smaller parts that consist of 30 minutes or less, you can significantly reduce the possibility of interruptions.
- Use automated retries To avoid disruptions to Amazon Batch jobs, set automated retries
  for jobs. Batch jobs can be disrupted for any of the following reasons: a non-zero exit code
  is returned, a service error occurs, or an instance reclamation occurs. You can set up to 10
  automated retries. For a start, we recommend that you set at least 1-3 automated retries. For
  information about tracking Amazon EC2 Spot interruptions, see Spot Interruption Dashboard.

For Amazon Batch, if you set the retry parameter, the job is placed at the front of the job queue. That is, the job is given priority. When you create the job definition or you submit the job in the Amazon CLI, you can configure a retry strategy. For more information, see submit-job.

```
$ aws batch submit-job --job-name MyJob \
    --job-queue MyJQ \
    --job-definition MyJD \
    --retry-strategy attempts=2
```

Use custom retries – You can configure a job retry strategy to a specific application exit code
or instance reclamation. In the following example, if the host causes the failure, the job can
be retried up to five times. However, if the job fails for a different reason, the job exits and the
status is set to FAILED.

```
"retryStrategy": {
    "attempts": 5,
    "evaluateOnExit":
    [{
        "onStatusReason" :"Host EC2*",
        "action": "RETRY"
```

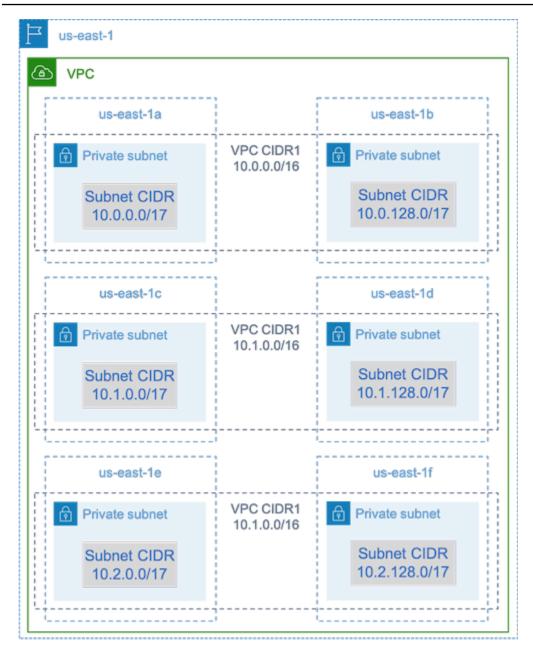
```
},{
    "onReason" : "*",
    "action": "EXIT"
}]
}
```

Use the Spot Interruption Dashboard – You can use the Spot Interruption Dashboard to track
Spot interruptions. The application provides metrics on Amazon EC2 Spot instances that are
reclaimed and which Availability Zones that Spot instances are in. For more information, see <a href="Spot Interruption Dashboard">Spot Interruption Dashboard</a>

# **Common errors and troubleshooting**

Errors in Amazon Batch often occur at the application level or are caused by instance configurations that don't meet your specific job requirements. Other issues include jobs getting stuck in the RUNNABLE status or compute environments getting stuck in an INVALID state. For more information about troubleshooting jobs getting stuck in RUNNABLE status, see <a href="Jobs stuck in a RUNNABLE status">Jobs stuck in a RUNNABLE status</a>. For information about troubleshooting compute environments in an INVALID state, see <a href="INVALID compute environment">INVALID compute environment</a>.

- Check Amazon EC2 Spot vCPU quotas Verify that your current service quotas meet the job requirements. For example, suppose that your current service quota is 256 vCPUs and the job requires 10,000 vCPUs. Then, the service quota doesn't meet the job requirement. For more information and troubleshooting instructions, see <a href="Mazon EC2 service quotas"><u>Amazon EC2 service quotas</u></a> and <a href="Mow do I"><u>How do I</u></a> increase the service quota of my Amazon EC2resources?
- Jobs fail before the application runs Some jobs might fail because of a
   DockerTimeoutError error or a CannotPullContainerError error. For troubleshooting
   information, see How do I resolve the "DockerTimeoutError" error in Amazon Batch?.
- Insufficient IP addresses The number of IP addresses in your VPC and subnets can limit the number of instances that you can create. Use Classless Inter-Domain Routings (CIDRs) to provide more IP addresses than are required to run your workloads. If necessary, you can also build a dedicated VPC with a large address space. For example, you can create a VPC with multiple CIDRs in 10.x.0.0/16 and a subnet in every Availability Zone with a CIDR of 10.x.y.0/17. In this example, x is between 1-4 and y is either 0 or 128. This configuration provides 36,000 IP addresses in every subnet.



- Verify that instances are registered with Amazon EC2 If you see your instances in the Amazon EC2 console, but no Amazon Elastic Container Service container instances in your Amazon ECS cluster, the Amazon ECS agent might not be installed on an Amazon Machine Image (AMI). The Amazon ECS Agent, the Amazon EC2 Data in your AMI, or the launch template might also not be configured correctly. To isolate the root cause, create a separate Amazon EC2 instance or connect to an existing instance using SSH. For more information, see <a href="Amazon ECS container">Amazon ECS container</a> agent configuration, Amazon ECS Log File Locations, and Compute resource AMIs.
- Review the Amazon Dashboard Review the Amazon Dashboard to verify that the expected job states and that the compute environment scales as expected. You can also review the job logs in CloudWatch.

Verify that your instance is created – If an instance is created, it means that your compute
environment scaled as expected. If your instances aren't created, find the associated subnets in
your compute environment to change. For more information, see <a href="Verify a scaling activity for an Auto Scaling group">Verify a scaling activity for an Auto Scaling group</a>.

We also recommend that you verify that your instances can fulfill your related job requirements. For example, a job might require 1 TiB of memory, but the compute environment uses a C5 instance type that's limited to 192 GB of memory.

- Verify that your instances are being requested by Amazon Batch Check Auto Scaling group history to verify that your instances are being requested by Amazon Batch. This is an indication of how Amazon EC2 tries to acquire instances. If you receive an error stating the Amazon EC2 Spot can't acquire an instance in a specific Availability Zone, this might be because the Availability Zone doesn't offer a specific instance family.
- Verify that instances register with Amazon ECS If you see instances in the Amazon EC2 console, but no Amazon ECS container instances in your Amazon ECS cluster, the Amazon ECS agent might not be installed on the Amazon Machine Image (AMI). Moreover, the Amazon ECS Agent, the Amazon EC2 Data in your AMI, or the launch template might not be configured correctly. To isolate the root cause, create a separate Amazon EC2 instance or connect to an existing instance using SSH. For more information, see CloudWatch agent configuration file: Logs section, Amazon ECS Log File Locations, and Compute resource AMIs.
- Open a support ticket If you're still experiencing issues after some troubleshooting and have
  a support plan, open a support ticket. In the support ticket, make sure to include information
  about the issue, workload specifics, the configuration, and test results. For more information, see
   Compare Amazon Web Services Support Plans.
- Review the Amazon Batch and HPC forums For more information, see the <u>Amazon Batch</u> and HPC forums.
- Review the Amazon Batch Runtime Monitoring Dashboard This dashboard uses a serverless
  architecture to capture events from Amazon ECS, Amazon Batch, and Amazon EC2 to provide
  insights into jobs and instances. For more information, see <a href="Amazon Batch Runtime Monitoring Dashboards Solution">Amazon Batch Runtime Monitoring Dashboards Solution</a>.

# **Troubleshooting Amazon Batch**

You might need to troubleshoot issues that are related to your compute environments, job queues, job definitions, or jobs. This chapter describes how to troubleshoot and resolve such issues in your Amazon Batch environment.

Amazon Batch uses IAM policies, roles, and permissions, and runs on Amazon EC2, Amazon ECS, Amazon Fargate, and Amazon Elastic Kubernetes Service infrastructure. To troubleshoot issues that are related to these services, see the following:

- <u>Troubleshooting IAM</u> in the *IAM User Guide*
- Amazon ECS troubleshooting in the Amazon Elastic Container Service Developer Guide
- Amazon EKS troubleshooting in the Amazon EKS User Guide
- Troubleshoot EC2 instances in the Amazon EC2 User Guide

#### **Contents**

- Amazon Batch
  - · INVALID compute environment
    - Incorrect role name or ARN
    - Repair an INVALID compute environment
  - Jobs stuck in a RUNNABLE status
  - Spot Instances not tagged on creation
  - Spot Instances not scaling down
    - Attach AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role in the Amazon Web Services Management Console
    - Attach AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role with the Amazon CLI
  - Can't retrieve Secrets Manager secrets
  - Can't override job definition resource requirements
  - Error message when you update the desiredvCpus setting
- Amazon Batch on Amazon EKS
  - INVALID compute environment
    - Unsupported Kubernetes version

- Instance profile doesn't exist
- Invalid Kubernetes namespace
- Deleted compute environment
- Nodes don't join the Amazon EKS cluster
- Amazon Batch on Amazon EKS job is stuck in RUNNABLE status
- Amazon Batch on Amazon EKS job is stuck in STARTING status
  - Scenario: Persisted Volume Claim Attach or Mount Failure
- Verify that the aws-auth ConfigMap is configured correctly
- RBAC permissions or bindings aren't configured properly

## Amazon Batch

Review the following topics to find review processes and potential solutions to common issues that you may encounter when using Amazon Batch.

#### **Topics**

- INVALID compute environment
- Jobs stuck in a RUNNABLE status
- Spot Instances not tagged on creation
- Spot Instances not scaling down
- Can't retrieve Secrets Manager secrets
- Can't override job definition resource requirements
- Error message when you update the desiredvCpus setting

## **INVALID** compute environment

It's possible that you might have incorrectly configured a managed compute environment. If you did, the compute environment enters an INVALID state and can't accept jobs for placement. The following sections describe the possible causes and how to troubleshoot based on the cause.



#### Important

Amazon Batch creates and manages multiple Amazon resources on your behalf and within your account, including Amazon EC2 Launch Templates, Amazon EC2 Auto Scaling

Amazon Batch 463

Groups, Amazon EC2 Spot Fleets, and Amazon ECS Clusters. These managed resources are configured specifically to ensure optimal Amazon Batch operation. Manually modifying these Batch-managed resources, unless explicitly stated in Amazon Batch documentation, may result in unexpected behavior resulting in INVALID Compute Environment, suboptimal instance scaling behavior, delayed workload processing, or unexpected costs. These manual modifications can not be deterministically supported by the Amazon Batch service. Always use the supported Batch APIs or the Batch console to manage your Compute Environments.

#### Incorrect role name or ARN

The most common cause for a compute environment to enter an INVALID state is that the Amazon Batch service role or the Amazon EC2 Spot Fleet role has an incorrect name or Amazon Resource Name (ARN). This is more common with compute environments that are created using the Amazon CLI or the Amazon SDKs. When you create a compute environment in the Amazon Web Services Management Console, Amazon Batch helps you choose the correct service or Spot Fleet roles. However, suppose that you manually enter the name or the ARN and enter them incorrectly. Then, the resulting compute environment is also INVALID.

However, suppose that you manually enter the name or ARN for an IAM resource in an Amazon CLI command or your SDK code. In this case, Amazon Batch can't validate the string. Instead, Amazon Batch must accept the bad value and attempt to create the environment. If Amazon Batch fails to create the environment, the environment moves to an INVALID state, and you see the following errors.

For an invalid service role:

```
CLIENT_ERROR - Not authorized to perform sts:AssumeRole (Service: AWSSecurityTokenService; Status Code: 403; Error Code: AccessDenied; Request ID: dc0e2d28-2e99-11e7-b372-7fcc6fb65fe7)
```

For an invalid Spot Fleet role:

```
CLIENT_ERROR - Parameter: SpotFleetRequestConfig.IamFleetRole is invalid. (Service: AmazonEC2; Status Code: 400; Error Code: InvalidSpotFleetRequestConfig; Request ID: 331205f0-5ae3-4cea-bac4-897769639f8d) Parameter: SpotFleetRequestConfig.IamFleetRole is invalid
```

One common cause for this issue is the following scenario. You only specify the name of an IAM role when using the Amazon CLI or the Amazon SDKs, instead of the full Amazon Resource Name (ARN). Depending on how you created the role, the ARN might contain a aws-service-role path prefix. For example, if you manually create the Amazon Batch service role using the procedures in Use service-linked roles for Amazon Batch, your service role ARN might look like the following.

```
arn:aws-cn:iam::123456789012:role/AWSBatchServiceRole
```

However, if you created the service role as part of the console first run wizard today, your service role ARN might look like the following.

```
arn:aws-cn:iam::123456789012:role/aws-service-role/AWSBatchServiceRole
```

This issue can also occur if you attach the Amazon Batch service-level policy (AWSBatchServiceRole) to a non-service role. For example, you may receive an error message that resembles the following in this scenario:

```
CLIENT_ERROR - User: arn:aws:sts::account_number:assumed-role/batch-replacement-role/aws-batch is not authorized to perform: action on resource ...
```

To resolve this issue, do one of the following.

- Use an empty string for the service role when you create the Amazon Batch compute environment.
- Specify the service role in the following format: arn:aws:iam::account\_number:role/aws-service-role/batch.amazonaws.com/AWSServiceRoleForBatch.

When you only specify the name of an IAM role when using the Amazon CLI or the Amazon SDKs, Amazon Batch assumes that your ARN doesn't use the aws-service-role path prefix. Because of this, we recommend that you specify the full ARN for your IAM roles when you create compute environments.

To repair a compute environment that's misconfigured this way, see Repair an INVALID compute environment.

## Repair an INVALID compute environment

When you have a compute environment in an INVALID state, update it to repair the invalid parameter. For an Incorrect role name or ARN, update the compute environment using the correct service role.

#### To repair a misconfigured compute environment

- 1. Open the Amazon Batch console at https://console.amazonaws.cn/batch/.
- 2. From the navigation bar, select the Amazon Web Services Region to use.
- 3. In the navigation pane, choose **Compute environments**.
- On the **Compute environments** page, select the radio button next to the compute environment to edit, and then choose Edit.
- On the **Update compute environment** page, for **Service role**, choose the IAM role to use with your compute environment. The Amazon Batch console only displays roles that have the correct trust relationship for compute environments.
- Choose **Save** to update your compute environment.

## Jobs stuck in a RUNNABLE status

Suppose that your compute environment contains compute resources, but your jobs don't progress beyond the RUNNABLE status. Then, it's likely that something is preventing the jobs from being placed on a compute resource and causing your job queues to be blocked. Here's how to know if your job is waiting for its turn or stuck and blocking the queue.

If Amazon Batch detects that you have a RUNNABLE job at the head and blocking the queue, you'll receive a Resource: Job gueue blocked events event from Amazon CloudWatch Events with the reason. The same reason is also updated into the statusReason field as a part of ListJobs and DescribeJobs API calls.

Optionally, you can configure the jobStateTimeLimitActions parameter through CreateJobQueue and UpdateJobQueue API actions.



#### (i) Note

Currently, the only action you can use with jobStateLimitActions.action is to cancel a job.

Jobs stuck in a RUNNABLE status 466

The jobStateTimeLimitActions parameter is used to specify a set of actions that Amazon Batch performs on jobs in a specific state. You can set a time threshold in seconds through the maxTimeSeconds field.

When a job has been in a RUNNABLE state with the defined statusReason, Amazon Batch performs the action specified after maxTimeSeconds have elapsed.

For example, you can set the jobStateTimeLimitActions parameter to wait up to 4 hours for any job in the RUNNABLE state that is waiting for sufficient capacity to become available. You can do this by setting statusReason to CAPACITY: INSUFFICIENT\_INSTANCE\_CAPACITY and maxTimeSeconds to 144000 before cancelling the job and allowing the next job to advance to the head of the job queue.

The following are the reasons that Amazon Batch provides when it detects that a job queue is blocked. This list provides the messages returned from the ListJobs and DescribeJobs API actions. These are also the same values you can define for the jobStateLimitActions.statusReason parameter.

- 1. **Reason:** All connected compute environments have insufficient capacity errors. When requested, Amazon Batch detects Amazon EC2 instances that experience insufficient capacity errors. Manually canceling the job will allow the subsequent job to move to the head of the queue but without resolving the service role issue(s), it is likely that the next job will also be blocked as well. It's best to manually investigate and resolve this issue.
  - statusReason message while the job is stuck:
     CAPACITY: INSUFFICIENT\_INSTANCE\_CAPACITY Service cannot fulfill the capacity requested for instance type [instanceTypeName]
  - reason used for jobStateTimeLimitActions: CAPACITY:INSUFFICIENT INSTANCE CAPACITY
  - **statusReason message after the job is canceled:** Canceled by JobStateTimeLimit action due to reason: CAPACITY:INSUFFICIENT\_INSTANCE\_CAPACITY

#### Note:

a. The Amazon Batch service role requires autoscaling: DescribeScalingActivities permission for this detection to work. If you use the <u>Service-linked role permissions</u> for Amazon Batch service-linked role (SLR) or the <u>Amazon managed policy</u>:
<u>AWSBatchServiceRole policy</u> managed policy, then you don't need to take any action because their permission policies are updated.

Jobs stuck in a RUNNABLE status 467

b. If you use the SLR or the managed policy, you must add the autoscaling:DescribeScalingActivities and ec2:DescribeSpotFleetRequestHistory permissions so that you can receive blocked job queue events and updated job status when in RUNNABLE. In addition, Amazon Batch needs these permissions to perform cancellation actions through the jobStateTimeLimitActions parameter even if they are configured on the job queue.

- c. In the case of a multi-node parallel (MNP) job, if the attached high-priority, Amazon EC2 compute environment experiences insufficient capacity errors, it blocks the queue even if a lower priority compute environment does experience this error.
- 2. **Reason:** All compute environments have a <a href="maxvCpus">maxvCpus</a> parameter that is smaller than the job requirements. Canceling the job, either manually or by setting the jobStateTimeLimitActions parameter on statusReason, allows the subsequent job to move to the head of the queue. Optionally, you can increase the maxvCpus parameter of the primary compute environment to meet the needs of the blocked job.
  - statusReason message while the job is stuck:

    MISCONFIGURATION: COMPUTE\_ENVIRONMENT\_MAX\_RESOURCE CE(s) associated with the job queue cannot meet the CPU requirement of the job.
  - reason used for jobStateTimeLimitActions:
     MISCONFIGURATION: COMPUTE\_ENVIRONMENT\_MAX\_RESOURCE
  - **statusReason message after the job is canceled:** Canceled by JobStateTimeLimit action due to reason: MISCONFIGURATION:COMPUTE\_ENVIRONMENT\_MAX\_RESOURCE
- 3. **Reason:** None of the compute environments have instances that meet the job requirements. When a job requests resources, Amazon Batch detects that no attached compute environment is able to accommodate the incoming job. Canceling the job, either manually or by setting the jobStateTimeLimitActions parameter on statusReason, allows the subsequent job to move to the head of the queue. Optionally, you can redefine the compute environment's allowed instance types to add the necessary job resources.
  - statusReason message while the job is stuck:
     MISCONFIGURATION: JOB\_RESOURCE\_REQUIREMENT The job resource
     requirement (vCPU/memory/GPU) is higher than that can be met by the
     CE(s) attached to the job queue.
  - reason used for jobStateTimeLimitActions:
     MISCONFIGURATION: JOB\_RESOURCE\_REQUIREMENT

- **statusReason message after the job is canceled:** Canceled by JobStateTimeLimit action due to reason: MISCONFIGURATION: JOB\_RESOURCE\_REQUIREMENT
- 4. Reason: All compute environments have service role issues. To resolve this, compare your service role permissions to the <u>Amazon managed policies for Amazon Batch</u> and address any gaps. Note:You can't configure a programmable action through the jobStateTimeLimitActions parameter to resolve this error.

It's a best practice to use the <u>Service-linked role permissions for Amazon Batch</u> to avoid similar errors.

Canceling the job, either manually or by setting the jobStateTimeLimitActions parameter on statusReason, allows the subsequent job to move to the head of the queue. Without resolving the service role issue(s), it is likely that the next job will also be blocked as well. It's best to manually investigate and resolve this issue.

- statusReason message while the job is stuck:
   MISCONFIGURATION:SERVICE\_ROLE\_PERMISSIONS Batch service role has a permission issue.
- 5. **Reason:** All compute environments are invalid. For more information, see <u>INVALID</u> <u>compute environment</u>. Note: You can't configure a programmable action through the jobStateTimeLimitActions parameter to resolve this error.
  - statusReason message while the job is stuck: ACTION\_REQUIRED CE(s) associated with the job queue are invalid.
- 6. **Reason:** Amazon Batch has detected a blocked queue, but is unable to determine the reason. Note: You can't configure a programmable action through the jobStateTimeLimitActions parameter to resolve this error. For more information about troubleshooting, see <a href="Why is my">Why is my</a> Amazon Batch job stuck in RUNNABLE on Amazon in re:Post.
  - statusReason message while the job is stuck: UNDETERMINED Batch job is blocked, root cause is undetermined.

In case you did not receive an event from CloudWatch Events or you received the unknown reason event, here are some common causes for this issue.

#### The awslogs log driver isn't configured on your compute resources

Amazon Batch jobs send their log information to CloudWatch Logs. To enable this, you must configure your compute resources to use the awslogs log driver. Suppose that you base

Jobs stuck in a RUNNABLE status 469

your compute resource AMI off of the Amazon ECS optimized AMI (or Amazon Linux). Then, this driver is registered by default with the ecs-init package. Now suppose that you use a different base AMI. Then, you must verify that the awslogs log driver is specified as an available log driver with the ECS\_AVAILABLE\_LOGGING\_DRIVERS environment variable when the Amazon ECS container agent is started. For more information, see <a href="Compute resource AMI">Compute resource AMI</a>. Specification and <a href="Tutorial">Tutorial</a>: Create a compute resource AMI.

#### Insufficient resources

If your job definitions specify more CPU or memory resources than your compute resources can allocate, then your jobs aren't ever placed. For example, suppose that your job specifies 4 GiB of memory, and your compute resources have less than that available. Then it's the case that the job can't be placed on those compute resources. In this case, you must reduce the specified memory in your job definition or add larger compute resources to your environment. Some memory is reserved for the Amazon ECS container agent and other critical system processes. For more information, see Compute resource memory management.

#### No internet access for compute resources

Compute resources need access to communicate with the Amazon ECS service endpoint. This can be through an interface VPC endpoint or through your compute resources having public IP addresses.

For more information about interface VPC endpoints, see <u>Amazon ECS Interface VPC Endpoints</u> (<u>Amazon PrivateLink</u>) in the <u>Amazon Elastic Container Service Developer Guide</u>.

If you do not have an interface VPC endpoint configured and your compute resources do not have public IP addresses, then they must use network address translation (NAT) to provide this access. For more information, see <a href="MAT gateways"><u>NAT gateways</u></a> in the *Amazon VPC User Guide*. For more information, see the section called "Create a VPC".

#### Amazon EC2 instance limit reached

The number of Amazon EC2 instances that your account can launch in an Amazon Web Services Region is determined by your EC2 instance quota. Certain instance types also have a perinstance-type quota. For more information about your account's Amazon EC2 instance quota including how to request a limit increase, see <a href="Mazon EC2 Service Limits"><u>Amazon EC2 Service Limits</u></a> in the *Amazon EC2 User Guide*.

#### Amazon ECS container agent isn't installed

The Amazon ECS container agent must be installed on the Amazon Machine Image (AMI) to let Amazon Batch run jobs. The Amazon ECS container agent is installed by default on Amazon ECS optimized AMIs. For more information about the Amazon ECS container agent, see <a href="Amazon ECS">Amazon ECS</a> container agent in the Amazon Elastic Container Service Developer Guide.

For more information, see Why is my Amazon Batch job stuck in RUNNABLE status? in re:Post.

## **Spot Instances not tagged on creation**

Spot Instance tagging for Amazon Batch compute resources is supported as of October 25, 2017. Before, the recommended IAM managed policy (AmazonEC2SpotFleetRole) for the Amazon EC2 Spot Fleet role didn't contain permissions to tag Spot Instances at launch. The new recommended IAM managed policy is called AmazonEC2SpotFleetTaggingRole. It supports tagging Spot Instances at launch.

To fix Spot Instance tagging on creation, follow the following procedure to apply the current recommended IAM managed policy to your Amazon EC2 Spot Fleet role. That way, any future Spot Instances that are created with that role have permissions to apply instance tags when they're created.

#### To apply the current IAM managed policy to your Amazon EC2 Spot Fleet role

- Open the IAM console at <a href="https://console.amazonaws.cn/iam/">https://console.amazonaws.cn/iam/</a>.
- 2. Choose **Roles**, and choose your Amazon EC2 Spot Fleet role.
- 3. Choose **Attach policy**.
- 4. Select the **AmazonEC2SpotFleetTaggingRole** and choose **Attach policy**.
- 5. Choose your Amazon EC2 Spot Fleet role again to remove the previous policy.
- 6. Select the **x** to the right of the **AmazonEC2SpotFleetRole** policy, and choose **Detach**.

## **Spot Instances not scaling down**

Amazon Batch introduced the **AWSServiceRoleForBatch** service-linked role on March 10, 2021. If no role is specified in the serviceRole parameter of the compute environment, this service-linked role is used as the service role. However, suppose that the service-linked role is used in an EC2 Spot compute environment, but the Spot role used doesn't include the

AmazonEC2SpotFleetTaggingRole managed policy. Then, the Spot Instance doesn't scale down. As a result, you will receive an error with the following message: "You are not authorized to perform this operation." Use the following steps to update the spot fleet role that you use in the spotIamFleetRole parameter. For more information, see <a href="Using service-linked roles">Using service-linked roles</a> and <a href="Creating a role to delegate permissions to an Amazon Service in the IAM User Guide">IAM User Guide</a>.

### **Topics**

- Attach AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role in the Amazon Web Services Management Console
- Attach AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role with the Amazon CLI

# Attach AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role in the Amazon Web Services Management Console

To apply the current IAM managed policy to your Amazon EC2 Spot Fleet role

- 1. Open the IAM console at <a href="https://console.amazonaws.cn/iam/">https://console.amazonaws.cn/iam/</a>.
- 2. Choose **Roles**, and choose your Amazon EC2 Spot Fleet role.
- 3. Choose Attach policy.
- 4. Select the AmazonEC2SpotFleetTaggingRole and choose Attach policy.
- 5. Choose your Amazon EC2 Spot Fleet role again to remove the previous policy.
- 6. Select the **x** to the right of the **AmazonEC2SpotFleetRole** policy, and choose **Detach**.

# Attach AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role with the Amazon CLI

The example commands assume that your Amazon EC2 Spot Fleet role is named AmazonEC2SpotFleetRole. If your role uses a different name, adjust the commands to match.

## To attach the AmazonEC2SpotFleetTaggingRole managed policy to your Spot Fleet role

 To attach the AmazonEC2SpotFleetTaggingRole managed IAM policy to your AmazonEC2SpotFleetRole role, run the following command using the Amazon CLI.

\$ aws iam attach-role-policy ackslash

```
--policy-arn arn:aws-cn:iam::aws:policy/service-role/
AmazonEC2SpotFleetTaggingRole \
    --role-name AmazonEC2SpotFleetRole
```

2. To detach the **AmazonEC2SpotFleetRole** managed IAM policy from your **AmazonEC2SpotFleetRole** role, run the following command using the Amazon CLI.

```
$ aws iam detach-role-policy \
    --policy-arn arn:aws-cn:iam::aws:policy/service-role/AmazonEC2SpotFleetRole \
    --role-name AmazonEC2SpotFleetRole
```

## **Can't retrieve Secrets Manager secrets**

If you use an AMI with an Amazon ECS agent that's earlier than version 1.16.0-1, then you must use the Amazon ECS agent configuration variable ECS\_ENABLE\_AWSLOGS\_EXECUTIONROLE\_OVERRIDE=true to use this feature. You can add it to the ./etc/ecs/ecs.config file to a new container instance when you create that instance. Or, you can add it to an existing instance. If you add it to an existing instance, you must restart the ECS agent after you add it. For more information, see <a href="Amazon ECS Container Agent Configuration">Amazon Elastic Container Service Developer Guide</a>.

## Can't override job definition resource requirements

The memory and vCPU overrides that are specified in the memory and vcpus members of the <u>containerOverrides</u> structure, which passed to <u>SubmitJob</u>, can't override the memory and vCPU requirements that are specified in the <u>resourceRequirements</u> structure in the job definition.

If you try to override these resource requirements, you might see the following error message:

"This value was submitted in a deprecated key and may conflict with the value provided by the job definition's resource requirements."

To correct this, specify the memory and vCPU requirements in the <u>resourceRequirements</u> member of the <u>containerOverrides</u>. For example, if your memory and vCPU overrides are specified in the following lines.

```
"containerOverrides": {
   "memory": 8192,
   "vcpus": 4
```

}

## Change them to the following:

Do the same change to the memory and vCPU requirements that are specified in the <u>containerProperties</u> object in the job definition. For example, if your memory and vCPU requirements are specified in the following lines.

```
{
  "containerProperties": {
    "memory": 4096,
    "vcpus": 2,
}
```

### Change them to the following:

## Error message when you update the desiredvCpus setting

You see the following error message when you use the Amazon Batch API to update the desired vCPUs (desiredvCpus) setting.

Manually scaling down compute environment is not supported. Disconnecting job queues from compute environment will cause it to scale-down to minvCpus.

This issue occurs if the updated desiredvCpus value is less than the current desiredvCpus value. When you update the desiredvCpus value, both of the following must be true:

- The desiredvCpus value must be between the minvCpus and maxvCpus values.
- The updated desiredvCpus value must be greater than or equal to the current desiredvCpus value.

## Amazon Batch on Amazon EKS

#### **Topics**

- INVALID compute environment
- Amazon Batch on Amazon EKS job is stuck in RUNNABLE status
- Amazon Batch on Amazon EKS job is stuck in STARTING status
- Verify that the aws-auth ConfigMap is configured correctly
- RBAC permissions or bindings aren't configured properly

Review the following topics to find review processes and potential solutions to common issues that you may encounter when using Amazon Batch on Amazon Elastic Kubernetes Service.

## **INVALID** compute environment

It's possible that you might have incorrectly configured a managed compute environment. If you did, the compute environment enters an INVALID state and can't accept jobs for placement. The following sections describe the possible causes and how to troubleshoot based on the cause.

## **Unsupported Kubernetes version**

You might see an error message that resembles the following when you use the CreateComputeEnvironment API operation or UpdateComputeEnvironmentAPI operation to create or update a compute environment. This issue occurs if you specify an unsupported Kubernetes version in EC2Configuration.

```
At least one imageKubernetesVersion in EC2Configuration is not supported.
```

To resolve this issue, delete the compute environment and then re-create it with a supported Kubernetes version.

You can perform a minor version upgrade on your Amazon EKS cluster. For example, you can upgrade the cluster from 1.xx to 1.yy even if the minor version isn't supported.

However, the compute environment status might change to INVALID after a major version update. For example, if you perform a major version upgrade from 1.xx to 2.yy. If the major version isn't supported by Amazon Batch, you see an error message that resembles the following.

```
reason=CLIENT_ERROR - ... EKS Cluster version [2.yy] is unsupported
```

To resolve this issue, specify a supported Kubernetes version when you use an API operation to create or update a compute environment.

Amazon Batch on Amazon EKS currently supports the following Kubernetes versions:

- 1.32
- 1.31
- 1.30
- 1.29
- 1.28
- 1.27
- 1.26
- 1.25

## Instance profile doesn't exist

If the specified instance profile does not exist, the Amazon Batch on Amazon EKS compute environment status is changed to INVALID. You see an error set in the statusReason parameter that resembles the following.

```
CLIENT_ERROR - Instance profile arn:aws-cn:iam::...:instance-profile/<name> does not
   exist
```

To resolve this issue, specify or create a working instance profile. For more information, see Amazon EKS node IAM role in the Amazon EKS User Guide.

## **Invalid Kubernetes namespace**

If Amazon Batch on Amazon EKS can't validate the namespace for the compute environment, the compute environment status is changed to INVALID. For example, this issue can occur if the namespace doesn't exist.

You see an error message set in the statusReason parameter that resembles the following.

```
CLIENT_ERROR - Unable to validate Kubernetes Namespace
```

This issue can occur if any of the following are true:

- The Kubernetes namespace string in the CreateComputeEnvironment call doesn't exist. For more information, see CreateComputeEnvironment.
- The required Role-Based Access Control (RBAC) permissions to manage the namespace are not configured correctly.
- Amazon Batch doesn't have access to the Amazon EKS Kubernetes API server endpoint.

To resolve this issue, see <u>Verify that the aws-auth ConfigMap is configured correctly</u>. For more information, see <u>Getting started with Amazon Batch on Amazon EKS</u>.

## **Deleted compute environment**

Suppose that you delete an Amazon EKS cluster before you delete the attached Amazon Batch on Amazon EKS compute environment. Then, the compute environment status is changed to

INVALID. In this scenario, the compute environment doesn't work properly if you re-create the Amazon EKS cluster with the same name.

To resolve this issue, delete and then re-create the Amazon Batch on Amazon EKS compute environment.

### Nodes don't join the Amazon EKS cluster

Amazon Batch on Amazon EKS scales down a compute environment if it determines that not all nodes joined the Amazon EKS cluster. When Amazon Batch on Amazon EKS scales down the compute environment, the compute environment status is changed to INVALID.



#### Note

Amazon Batch doesn't change the compute environment status immediately so that you can debug the issue.

You see an error message set in the statusReason parameter that resembles ones of the following:

Your compute environment has been INVALIDATED and scaled down because none of the instances joined the underlying ECS Cluster. Common issues preventing instances joining are the following: VPC/Subnet configuration preventing communication to ECS, incorrect Instance Profile policy preventing authorization to ECS, or customized AMI or LaunchTemplate configurations affecting ECS agent.

Your compute environment has been INVALIDATED and scaled down because none of the nodes joined the underlying Amazon EKS Cluster. Common issues preventing nodes joining are the following: networking configuration preventing communication to Amazon EKS Cluster, incorrect Amazon EKS Instance Profile or Kubernetes RBAC policy preventing authorization to Amazon EKS Cluster, customized AMI or LaunchTemplate configurations affecting Amazon EKS/Kubernetes node bootstrap.

When using a default Amazon EKS AMI, the most common causes of this issue are the following:

• The instance role isn't configured correctly. For more information, see Amazon EKS node IAM role in the Amazon EKS User Guide.

• The subnets aren't configured correctly. For more information, see Amazon EKS VPC and subnet requirements and considerations in the Amazon EKS User Guide.

• The security group isn't configured correctly. For more information, see Amazon EKS security group requirements and considerations in the Amazon EKS User Guide.



#### Note

You may also see an error notification in the Personal Health Dashboard (PHD).

## Amazon Batch on Amazon EKS job is stuck in RUNNABLE status

An aws-auth ConfigMap is automatically created and applied to your cluster when you create a managed node group or a node group using eksctl. An aws-auth ConfigMap is initially created to allow nodes to join your cluster. However, you also use the aws-authConfigMap to add rolebased access control (RBAC) access to users and roles.

To verify that the aws-auth ConfigMap is configured correctly:

Retrieve the mapped roles in the aws-auth ConfigMap:

```
$ kubectl get configmap -n kube-system aws-auth -o yaml
```

Verify that the roleARN is configured as follows.

rolearn: arn:aws-cn:iam::aws\_account\_number:role/AWSServiceRoleForBatch



### Note

You can also review the Amazon EKS control plane logs. For more information, see Amazon EKS control plane logging in the Amazon EKS User Guide.

To resolve an issue where a job is stuck in a RUNNABLE status, we recommend that you use kubect1 to re-apply the manifest. For more information, see Step 2: Prepare your Amazon EKS cluster for Amazon Batch. Or, you can use kubectl to manually edit the aws-auth ConfigMap. For more information, see Enabling IAM user and role access to your cluster in the Amazon EKS User Guide.

## Amazon Batch on Amazon EKS job is stuck in STARTING status

A Job may remain in STARTING status when the Pod is stuck in PENDING on ContainerCreating for any long running requests from the kubelet (pull, log, exec, and attach) until the Pod startup issue is resolved or the Job is terminated. In the qualifying scenarios below Amazon Batch will terminate the job on your behalf, otherwise the job must be terminated manually using the TerminateJob API.

To verify the reason a Job may be stuck in STARTING, use <u>Tutorial: Map a running job to a pod and</u> a node to find the podName, and describe the Pod:

```
% kubectl describe pod aws-batch.000c8190-87df-31e7-8819-176fe017a24a -n my-aws-batch-
namespace
Name:
                  aws-batch.000c8190-87df-31e7-8819-176fe017a24a
Namespace:
                  my-aws-batch-namespace
Containers:
  default:
    State:
                    Waiting
      Reason:
                    ContainerCreating
    Ready:
                    False
Conditions:
  Type
                               Status
  PodReadyToStartContainers
                               False
  Initialized
                               True
                               False
  Ready
  ContainersReady
                               False
  PodScheduled
                               True
Events:
  Type
           Reason
                        Age
                                From
                                         Message
  Warning FailedMount
                        2m32s
                                kubelet
                                         Unable to attach or mount volumes: ...
```

Consider configuring your EKS cluster to <u>Send control plane logs to CloudWatch Logs</u> for full visibility.

#### Scenario: Persisted Volume Claim Attach or Mount Failure

Jobs using Persistent Volume Claims where the volume fails to attach or mount are candidates for termination. This can be a result of an incorrectly configured Job Definition. See <u>Tutorial: Create a single-node job definition on Amazon EKS resources</u> for more details.

## Verify that the aws-auth ConfigMap is configured correctly

To verify that the aws-auth ConfigMap is configured correctly:

1. Retrieve the mapped roles in the aws-auth ConfigMap.

```
$ kubectl get configmap -n kube-system aws-auth -o yaml
```

Verify that the roleARN is configured as follows.

rolearn: arn:aws-cn:iam::aws\_account\_number:role/AWSServiceRoleForBatch



The path aws-service-role/batch.amazonaws.com/ has been removed from the ARN of the service-linked role. This is because of an issue with the aws-auth configuration map. For more information, see Roles with paths do not work when the path is included in their ARN in the aws-authconfigmap.

## Note

You can also review the Amazon EKS control plane logs. For more information, see Amazon EKS control plane logging in the Amazon EKS User Guide.

To resolve an issue where a job is stuck in a RUNNABLE status, we recommend that you use kubectl to re-apply the manifest. For more information, see <a href="Step 2">Step 2</a>: Prepare your Amazon EKS cluster for Amazon Batch. Or, you can use kubectl to manually edit the aws-auth ConfigMap. For more information, see <a href="Enabling IAM user and role access to your cluster">Enabling IAM user and role access to your cluster</a> in the Amazon EKS User Guide.

# RBAC permissions or bindings aren't configured properly

If you experience any RBAC permissions or binding issues, verify that the aws-batch Kubernetes role can access the Kubernetes namespace:

```
$ kubectl get namespace namespace --as=aws-batch
$ kubectl auth can-i get ns --as=aws-batch
```

You can also use the **kubect1 describe** command to view the authorizations for a cluster role or Kubernetes namespace.

```
$ kubectl describe clusterrole aws-batch-cluster-role
```

The following is example output.

Name:	aws-batch-cluster-role		
Labels:	<none></none>		
Annotations:	<none></none>		
PolicyRule:			
Resources		Non-Resource URLs	Resource Names
Verbs			
		F.3	F.3
configmaps	L-1-7	[]	[]
[get list wa nodes	tcnj	F7	r1
[get list wa	+ch]	[]	[]
pods	tenj	[]	[]
[get list wa	tch]	LJ	LJ
daemonsets.		[]	[]
[get list wa		L.3	L3
deployments		[]	[]
[get list wa			
replicasets		[]	[]
[get list wa	tch]		
statefulset		[]	[]
[get list wa	tch]		
clusterrole	bindings.rbac.authorization.k8s.io	[]	[]
[get list]			

<pre>clusterroles.rbac.authorization.k8s.io [get list]</pre>	[]	[]	
namespaces	[]	[]	
[get] events	[]	[]	
[list]			

\$ kubectl describe role aws-batch-compute-environment-role -n my-aws-batch-namespace

The following is example output.

Name: aws-batch-compute-environme Labels: <none> Annotations: <none></none></none>	ent-role		
PolicyRule: Resources	Non-Resource URLs	Resource Names	Verbs
<pre>pods get list watch delete patch]</pre>	[]	[]	[create
serviceaccounts rolebindings.rbac.authorization.k8s.io roles.rbac.authorization.k8s.io	[] []	[] []	<pre>[get list] [get list] [get list]</pre>

To resolve this issue, re-apply the RBAC permissions and rolebinding commands. For more information, see <a href="Step 2">Step 2</a>: Prepare your Amazon EKS cluster for Amazon Batch.

# **Resource: Amazon Batch service quotas**

The following table provides the service quotas for Amazon Batch that can't be changed. Each quota is Region specific.

Resource	Quota
Maximum number of job queues. For more information, see <u>Job queues</u> .	50
Maximum number of compute environments across Amazon ECS and Amazon EKS. For more information, see <i>Compute environments for Amazon Batch</i> .	50
Maximum number of compute environments per Amazon EKS cluster.	5
Maximum number of compute environments for each job queue	3
Maximum number of job dependencies for a job	20
Maximum job definition size (for <a href="RegisterJobDefinition">RegisterJobDefinition</a> API operations)	24 KiB
Maximum job payload size (for <a href="SubmitJob">SubmitJob</a> API operations)	30 KiB
Maximum array size for array jobs	10000
Maximum number of jobs in SUBMITTED state	1000000
Maximum number of transactions per second (TPS) for each account for <a href="SubmitJob">SubmitJob</a> operations	50
Maximum number of consumable resources	50k

Depending on how you use Amazon Batch, additional quotas might apply. To learn about Amazon EC2 quotas, see <a href="Manazon EC2 Service Quotas"><u>Amazon EC2 Service Quotas</u></a> in the *Amazon Web Services General Reference*. For more information about Amazon ECS quotas, see <a href="Amazon ECS Service Quotas"><u>Amazon Web Services General Reference</u></a>. For more information about Amazon EKS quotas, see <a href="Amazon EKS"><u>Amazon EKS</u></a> Service Quotas in the *Amazon Web Services General Reference*.

# **Document history**

The following table describes the important changes to the documentation since the initial release of Amazon Batch. We also update the documentation frequently to address the feedback that you send us.

Change	Description	Date
Adds support for EKS AL 2023 AMIs	How to upgrade from EKS AL2 to EKS AL2023	June 24, 2025
Adds support for FireLens and ECS Exec command	Adds support for FireLens and ECS Exec command.	April 15, 2025
Adds support for resource- aware scheduling for Amazon Batch	Adds support for resource- aware scheduling for Amazon Batch for Amazon Elastic Container Service, Amazon Elastic Kubernetes Service, and Amazon Fargate.	February 27, 2025
Updated Amazon Batch supported Amazon EKS versions	Updated the Amazon EKS versions that Amazon Batch supports to remove version 1.22.	March 11, 2024
Updated Amazon Batch supported Amazon EKS versions	Updated the Amazon EKS versions that Amazon Batch supports to include version 1.29.	February 29, 2024
Automated job retries	Corrected the code sample.	February 29, 2024
Adds support for multi-con tainer jobs for Amazon Batch	Adds support for multi- container jobs for Amazon Batch for Amazon Elastic Container Service, Amazon	February 28, 2024

Elastic Kubernetes Service, and Amazon Fargate. **Updated Amazon Batch** Updated the Amazon EKS January 27, 2024 versions that Amazon Batch supported Amazon EKS versions supports to include version 1.28 Updated BatchServiceRolePo December 5, 2023 **BatchServiceRolePolicy** licy and AWSBatchServiceRole Updated to add support for describing Spot Fleet request history and Amazon EC2 Auto Scaling activities. **AWSBatchServiceRole** Updated to add statement IDs, grant Amazon Batch permissions to ec2:Descr ibeSpotFleetReques tHistory and autoscaling:Descri beScalingActivitie s. Amazon Batch on Amazon Amazon Batch adds support October 25, 2022 **EKS** for running jobs on Amazon EKS clusters. June 6, 2022 Cross-service confused deputy Amazon Batch now provides a workaround for the confused prevention for Amazon Batch deputy security issue, which arises when an entity (a

service or an account) is

to perform an action.

coerced by a different entity

Added support for configuri ng interface VPC endpoints powered by Amazon PrivateLi nk. This means you can create a private connection between your VPC and Amazon Batch without requiring access through a NAT instance, a VPN connection, or Amazon Direct Connect.	April 15, 2022
Amazon Batch enhanced support updates to compute environments.	April 14, 2022
Amazon Batch updated existing managed policies.	December 6, 2021
Amazon Batch adds support for adding scheduling policies to job queues.	November 9, 2021
Amazon Batch adds support for adding Amazon EFS file systems to your job definitions.	April 1, 2021
Amazon Batch adds the <b>AWSServiceRoleForBatch</b> service-linked role.	March 10, 2021
Amazon Batch adds support for running jobs on Fargate resources.	December 3, 2020
	ng interface VPC endpoints powered by Amazon PrivateLi nk. This means you can create a private connection between your VPC and Amazon Batch without requiring access through a NAT instance, a VPN connection, or Amazon Direct Connect.  Amazon Batch enhanced support updates to compute environments.  Amazon Batch updated existing managed policies.  Amazon Batch adds support for adding scheduling policies to job queues.  Amazon Batch adds support for adding Amazon EFS file systems to your job definitio ns.  Amazon Batch adds the AWSServiceRoleForBatch service-linked role.  Amazon Batch adds support for running jobs on Fargate

Resource tagging	Amazon Batch adds support for adding metadata tags to your compute environments, job definitions, job queues, and jobs.	October 7, 2020
Secrets	Amazon Batch adds support for passing secrets to jobs.	October 1, 2020
Logging	Amazon Batch adds support for specifying additional log drivers for jobs.	October 1, 2020
Allocation strategies	Amazon Batch adds support for multiple strategies to choose instance types.	October 16, 2019
EFA support	Amazon Batch adds support for Elastic Fabric Adapter (EFA) devices.	August 2, 2019
GPU scheduling	Amazon Batch adds GPU scheduling. With this feature, you can specify the number of GPUs each job requires, and Amazon Batch scales up instances accordingly.	April 4, 2019
Multi-node parallel jobs	Amazon Batch adds support for multi-node parallel jobs. You can use this feature run single jobs that span over multiple Amazon EC2 instances.	November 19, 2018
Resource-level permissions	Amazon Batch supports resource-level permissions on several API operations.	November 12, 2018

Amazon EC2 Launch template support	Amazon Batch adds support for using launch templates with compute environments.	November 12, 2018
Amazon Batch job timeouts	Amazon Batch adds support for job timeout. With this support, you can configure a specific timeout duration for your jobs so that if a job runs longer than they should, Amazon Batch terminates the job.	April 5, 2018
Amazon Batch jobs as EventBridge targets	Amazon Batch jobs are made available as EventBridge targets. By creating simple rules, you can match events and submit Amazon Batch jobs in response to them.	March 1, 2018
CloudTrail auditing for Amazon Batch	CloudTrail can audit calls made to Amazon Batch API actions.	January 10, 2018
Array jobs	Amazon Batch adds support for array jobs. You can use array jobs for parameter sweep and Monte Carlo workloads.	November 28, 2017
Expanded Amazon Batch tagging	Amazon Batch expands support for the tagging function. You can use this function to specify tags for Amazon EC2 Spot Instances launched within managed compute environments.	October 26, 2017

Amazon B	atch	event	stream
for EventB	rida	e	

Amazon Batch adds the event stream for EventBridge. You can use Amazon Batch event stream to receive near realtime notifications regarding the state of jobs that are submitted to your job queues. October 24, 2017

## Automated job retries

Amazon Batch adds support for job retries. With this update, you can apply a retry strategy to your jobs and job definitions that allows your jobs to be automatically retried if they fail. March 28, 2017

# Amazon Batch general availability

Amazon Batch is introduce d, designed as a means for you to run batch computing workloads on the Amazon Web Services Cloud. January 5, 2017